Blender Wiki PDF Manual

Table of Contents

About Blender Wiki PDF Manual **1**
- Modifications to the wiki content 1
- Previous versions 1
- Where to find this PDF 2
 - Alternatives 2
- How to contact me 2

Open Content License **3**
- Introduction 4
- Starting 4
- Interaction in 3D 5
- Data System 5
- Modeling 6
- Modifiers and Deformation 7
- Lighting 7
- Materials 8
- Textures 9
- World and Ambient Effects 10
- Animation Basics 10
- Armatures and Character Animation 10
- Constraints 11
- Advanced Animation 12
- Effects and Physical Simulation 12
- Rendering 12
- Compositing with nodes 13
- Editing Sequences 13
- Extending Blender 14
- Game Engine 14

Introduction **16**

What is Blender? **17**

Blenderâ ← s History **18**
- Version/Revision Milestones 19

About Free Software and the GPL **21**

Getting support: the Blender community **22**

Who uses Blender? **23**

About this Manual **24**

Audience **25**

Learning CG and Blender **26**

Installing the Binaries **28**
- Downloading and installing the binary distribution 28

Table of Contents

Hardware Support **29**
- Developers platforms 29

Compiling the Source **31**

Compiling the Plugins **32**

Python, the Scripting Language **33**

Windows **35**
- Quick Install 35
- In-depth Instructions 35
- Portable Install 36

Linux **37**
- Quick Install 37
- In-depth Instructions 37
 - How to add Blender's icons in KDE 37
 - How to add Blender's icons in GNOME 38

OSX **46**

Other supported operating systems **47**
- FreeBSD 47
- Irix 47
- Solaris 47

Other unsupported operating systems **48**
- MorphOS 48

Configure your Blender **49**

Setting up a user directory structure **50**
- Folder explanations 50

Internet Access **52**

Introduction **53**

Blender's Interface Concept **54**

About Interface Internationalization **55**

Keyboard and mouse **56**
- Conventions in this Manual 56
- General Usage 56
- Mouse Button Emulation 56
- NumPad Emulation 57

The Window System **58**
- The default scene 58
 - The main top menu 59

Table of Contents

The Window System
- 3D Window View....59
- 3D Window Header....59
- Buttons Window Header....60
- Buttons Window View....60

Changing Window Frames....61
- Maximizing a window....61
- Splitting a window....61
- Joining two windows....61
- Changing window size....62

Application Frame....63

The Window Header....64
- Hiding a header....64
- Showing a header....65
- Header position....65
- Window type button....66
- Menus and buttons....66

The Console Window....67
- Windows 2000/Xp/Vista....67
- Linux....68
- Console Window Status & Error Messages....69

Window types....71
- See also....72

Screens....73

Screens Configuration....75
- Adding a new Screen....75
- Deleting a Screen....75
- Rearranging a Screen....75
- Overriding Defaults....75
- Additional Layouts....75

Scenes....77

Scenes configuration....78
- Adding a new Scene....78
- Deleting a Scene....78

Configuration....79

Contexts....80

Menus....81

Table of Contents

Panels **82**

Buttons and Controls **83**
- Operation Button 83
- Toggle Button 83
- Radio Buttons 83
- Number Buttons 83
- Menu Buttons 84
- Color Selector controls 84
- Cascade Buttons 85

Your First Animation in 30 plus 30 Minutes Part I **86**
- Warming up 86
- Building the body 88
 - Mirror modelling 89
 - Arms and Legs 91
 - The Head 92
 - SubSurfaces (subsurf) 93
 - Constrained Scaling 94
- Let's see what Gus looks like 94
 - Camera setup 95
 - The Ground 96
 - Lights 97
 - Rendering 98
 - Saving our work 99
- Materials and Textures 100
 - Eyes and detail 102
 - Mouth 103
 - Eyes material 105
 - Rendering 106
 - Saving 106

Your First Animation in 30 plus 30 Minutes Part II **108**
- Rigging 108
- Skinning 110
 - Vertex groups 112
- Posing 113
 - Original position 113
 - Inverse Kinematics 113
 - Forward Kinematics 113
- Gus walks! 116

Opening Files **118**
- Description 118
- Options 118
 - Loading the UI 118
 - Navigating your Hard Disk 118
 - Other File Open Options 119

Saving Files **120**
- Description 120
- Options 120

Table of Contents

Saving Files
- Compress Files ... 120
- Hints ... 121

Other File Menu Options ... **122**

Rendering ... **123**
- Description ... 123
- Options ... 123
- Saving to disk ... 124
- Hints ... 124

Screenshots ... **125**

Setting the default scene ... **126**

User Preferences and Themes ... **127**
- View & Controls ... 128
- Edit Methods ... 128
- Language & Fonts ... 129
- Themes ... 129
 - Customizing Theme ... 129
 - Creating / Saving a Theme ... 129
 - Customizing Icons ... 129
- Auto Save ... 130
- System & OpenGL ... 131
- File Paths ... 132

Saving your Preferences ... **133**

Description ... **134**

Getting Started ... **135**

Undo ... **136**

Redo ... **137**

History ... **138**

Help system ... **139**
- About Blender ... 139
- Blender/Python Scripting API ... 139
- Getting Started ... 139
- HotKey and MouseAction Reference ... 140
- ID Property Browser ... 140
- Manual ... 140
- Python Scripting Reference ... 140
- Release Notes ... 140
- Scripts Help Browser ... 140
- Tutorials ... 140
- Websites ... 141

Table of Contents

Help system
System....141

Introduction....142

Introduction....143

Rotating the View....144
Description....144
Options....144
TrackBall/Turntable....145

Panning the View....146
Description....146

Zooming the View....147
Description....147

Fly mode....148
Description....148
Options....148

Perspective and Orthographic Projection....149
Description....149
Options....149
Technical Details....150
Perspective definition....150
Orthographic definition....150

View Shading....151
Description....151

Local View....152
Description....152

View Clipping Border....153
Description....153
Examples....153

Introduction....154

3D Window Header....155
View Menu....155
Select Menu....157
Object Menu....159
Mode List....160
ViewPort Shading List....161
Rotation Selector....161
Transform (Manipulator) Selectors....161
Layer Selector....162
Render Button....162

Table of Contents

Using the 3D Window **163**
- Description 163
- Edit Mode 163
- 3D Window Toolbox (Popup Menu) 163
 - Add 163
 - Edit 165
 - Select 165
 - Transform 165
 - Object 166
 - View 166
 - Render 166

The Camera **168**
- Add a new camera 168
- Change active camera 168
- Move active camera to view 168
- Camera Settings 169
- Camera Navigation 170
 - Roll, Pitch, Dolly, and Track 170
 - Aiming the camera in Flymode 170

Layers **171**
- Description 171
- Options 171
 - Viewing layers 171
 - Locking the View 172
 - Multiple Layers 172

Moving objects between layers **173**

Animating Layers **174**

Example of Object layer arrangement **175**

Layer Naming Script **176**

Local or Global View **177**
- Description 177
- Examples 177

Introduction **178**

What is Grease Pencil? **179**

Preparing **180**

Drawing **181**
- Quick Usage (For Just A Few Strokes) 181
- Easier usage (for drawing more complex sketches) 181
- Special Tricks in â← Draw Modeâ← 181
- For Tablet Users 181
- Sensitivity When Drawing 181

Table of Contents

Drawing
- Additional notes ... 182
- Drawing Planes ... 182

Layers ... **183**

Animation of the Sketches ... **184**
- Onion Skinning ... 184
- Adjusting Timing of Sketches ... 184
- Copying Sketches ... 185

Converting Sketches to Other Forms ... **186**

Introduction ... **187**
- Menu Options ... 187

Grab/Move ... **188**
- Description ... 188
- Options ... 188

Rotate ... **189**
- Description ... 189
- Options ... 189

Scale ... **190**
- Description ... 190
- Options ... 190

Precision ... **191**
- Description ... 191
- Options ... 191

Numeric transformations ... **192**
- Description ... 192

Clear transformations ... **193**
- Description ... 193
- Options ... 193

Align to ... **194**
- Description ... 194

Manipulators ... **195**
- Manipulator types ... 195
- Other Manipulator controls ... 196
- Manipulator Preferences ... 196
- Transform Orientation ... 197

Gestures ... **200**
- Scale ... 200
- Rotate ... 200
- Translate ... 200

Table of Contents

Normal axis locking **202**
- Move/Translate 202
- Rotate 202
- Scale 203

Plane axis locking **204**
- Move/Translate 204
- Rotate 204
- Scale 204

Pivot **206**
- Active Object 206
 - In Object mode 206
 - In Edit mode 206
- Individual Center 208
 - In Object mode 208
 - In Edit mode 209
- 3D Cursor 210
 - Positioning the 3D cursor 210
 - Transformation 211
- Median Point 212
 - In Object mode 212
 - In Edit mode 212
 - Transformation 212
- Bounding Box Center 212
 - In Object mode 213
 - In Edit mode 213

Proportional Edit **214**
- Options 214
- Examples 216

Transform Orientations **217**
- Axis Locking 217
- Plane Locking 217
- Custom Orientations 217
- Transform Orientations Panel 218

Transform Properties Panel **219**
- Options in Object mode 219
 - Transform Properties Locking 220

Overview **221**

Outliner and OOPS Schematic **222**

Users (Sharing) **223**
- Fake User 223

Copying and Linking Objects Between Scenes **224**

Table of Contents

Appending or Linking Across Files **225**
Proxy Objects 225

Pack and Unpack Data **226**
Unpack Data 226

Datablocks **227**
Scene Datablock 227
Object Datablock 227
ObData Datablocks 227
Curve Datablock 227
Camera Datablock 227
Lattice Datablock 227
Lamp Datablock 227
Metaball Datablock 227
Mesh Datablock 228
Material Datablock 228
Texture Datablock 228

Copying and Linking Datablocks **229**
Copying and Linking Scene Datablock 229
Copying and Linking Object Datablocks 229
Copying and Linking other Datablocks 229
Unlinked Datablocks 229

Scene Management Structure **230**

Planning Your Timeline **231**

Current Screen Layout and Scene **232**
Loading the UI with â← Fileâ← â â← Openâ← 232

Working with Scenes **233**
Adding a Scene 233
Naming a Scene 234
Linking to a Scene 234
Removing a scene from the file 234

Outliner window **235**
Description 235
Outliner view 235
Oops Schematic view 236

Selecting the Outliner Window Type **237**

Using the Outliner view **238**
Selecting and activating 238
Toggling object-level restrictions 239
Searching 240
Filtering the display 240
Example 241

Table of Contents

Using the Oops Schematic **242**
- Layout of the Oops Schematic 242
- Making sense of the Oops Schematic 242
- The Oops Schematic Header 243

Linked Libraries Overview **245**
- General Procedure 245
 - Folder and File Organization 245
 - Appending library objects into your current project 245
- Reusing Objects (Meshes, Curves, Cameras, Lights, â ←¦) 246
- Reusing Material/Texture Settings 247
- Reusing Node Layouts 248
- Proxy Objects 248

Modeling in Blender **250**

Object Mode **251**
- Erase 251
- Join 251
- Select Links 251

Introduction **253**

Selections and the Active Object **254**

Point Selection **255**

Rectangular or Border Select **256**
- Description 256
- Example 256
- Hints 256

Lasso Select **257**
- Description 257
- Usage 257

Menu Selection **258**
- Select Grouped 258
 - Description 258
 - Options 258
- Select linked 259
 - Description 259
 - Options 259
- Select All by Type 260
 - Description 260
 - Options 260
- Select All by Layer 260
 - Description 261
 - Options 261
- Other Menu Options 261

Table of Contents

Moving (translating) objects **263**
Moving Object(s) Individually 263
Description 263
Options 263
Hints 266
Moving/Translating Object(s) by Changing Attributes 266
Description 266
Options 266

Rotating objects **267**
Description 267
Options 267
Axis of rotation Constraint 267
Point of rotation 267
Hints 268

Scaling objects **269**
Description 269
Options 269
Axis of scale Constraint 269
Center point of scale 269
Mirroring objects 269

Skinning and Cloning Objects **271**
Options 271

Regrouping Objects **272**
Parenting objects 272
Description 272
Options 273
Parenting Example 273
Hints 274
Separating Objects 274
Description 274
Options 275
Grouping objects 275
Description 275
Options 275
Select Grouped 277
Description 277
Options 277
Example 278

Tracking **279**
Description 279
Options 279
TrackTo Constraint 279
LockTrack Constraint 280
Old Track 282
Hints 282
Invalid Tracking or settings 283

Table of Contents

Duplicate **284**
Description 284
Examples 284

Linked Duplicates **285**
Description 285
Examples 285

Procedural Duplication **286**

Hints **287**

DupliVerts **288**
DupliVerts as an Arranging Tool 288
DupliVerts to Model a Single Object 293
See also 294

DupliFaces **295**

DupliGroup **297**
Basic Usage 297
DupliGroup and Dynamic Linking 297
Example 297
Making a DupliGroup Objects Real 298
Armatures 298

DupliFrames **299**
Modeling using DupliFrames 299
Arranging objects with DupliFrames 303
More Animation and Modeling 304

Edit Mode **308**
How to enter Edit mode 308
Visualization 308
Mesh panels 309

Basic Editing **310**
Mirror Axis and Modifier 310
Specials 311
Mesh Undo/Redo 312

Vertices, Edges and Faces **314**
Vertices 314
Edges 314
Faces 314

Edge Loops and Face Loops **316**
Edge Loops 316
Face Loops 317

Table of Contents

Basic Mesh Objects **318**
Description 318
Options 318

Mesh smoothing **322**
Smoothing the entire mesh 322
Manually smoothing parts of a mesh 323
Auto Smooth 323

Smoothing the mesh geometry **325**

Choosing which elements are selectable **326**
Select Mode popup 326
Select Mode header widgets 326
Selected elements after switching selection mode 327

Basic selection **328**
Adding to a selection 328
Selecting different types of faces 328

Selecting elements in a Region **329**
Rectangular region (Border select) 329
Circular region 329
Lasso region 330

Basic Mesh Modeling **331**
Extrude 331
The Blade 331
The Handle 334
The Hilt 334
Mirror 335
In Edit Mode 336
In Object Mode 338
In Conclusion 339
Spin and SpinDup 339
Spin 339
Spin Duplicate 342
Screw 345
Warp Tool 346
To Sphere 348

Advanced Mesh Modeling **350**
Symmetrical Modeling 350
Noise 352
Example 353

Booleans Operations **355**

Boolean Tools **356**
Description 356
Options 356

Table of Contents

Boolean Modifier **358**
Description 358
Options 359

Examples **360**
Intersect 360
Union 360
Difference 360
More complex examples with Boolean modifiers 361
Multiple Modifiers 361
Animation 361

Technical Details **363**

Limitations & Workarounds **364**

Edge Tools **365**
Introduction 365
Edge Selection 366
Description 366
Options 366
Edge Loop Selection 367
Description 367
Examples 367
Technical Details 367
Edge Ring Selection 367
Description 368
Examples 368
Technical Details 368
Rotate Edge CW / Rotate Edge CCW 368
Description 368
Examples 368
Edge Slide 369
Description 369
Options 369
Examples 369
Limitations & Workarounds 370
Bevel 370
Description 370
Options 371
Hints 372
Bevel Selected Edges 372
Examples 372

Face Tools **373**
Face Selection 373
Description 373
Options 374
Face Loop Selection 374
Examples 374
Technical Details 374

Table of Contents

Advanced Edge and Face Tools **376**
Loop to Region and Region to Loop 376
Description 376
Examples: Loop to Region 376
Example: Region to Loop 377
Loop Subdivide 377
Description 377
Options 378
Examples 379
Delete Edge Loop 380
Description 381
Limitations & Workarounds 381
Examples 381
Knife Subdivide 382
Description 382
Options 382
Examples 383
Limitations & Workarounds 384
Optimizations 385

Snap to Mesh **386**
Description 386
Options 386
Snap Mode 386
Snap Element 386
Align Rotation 387

Vertex Groups **388**
Why use Vertex Groups? 388
Vertex Groups Basics 388
Creating a Vertex Group 389
Naming Vertex Groups 389
Assigning Vertices to a Group 390
Seeing a Vertex Group 390
Removing Vertices from a Group 391
Deselecting Vertices 391
Deleting a Group 391
Using Vertex Groups in Practice 392
Duplicating Parts 392
Simplifying a Vertex Group 392
Combining Groups 393
Focus on a part of your model 393
Separating a part into its own 393
Finding Groups for a Vertex 393
About Weight 394

Weight paint **395**
Paint-Panel 396
Tools 397
Scripts in Paint-Menu 399
Weight-Painting for Bones 399
Weight-Painting for Particles 400

Table of Contents

Subdivision Surfaces **401**
- Description 401
- Options 401
- Hints 402
- Examples 402
- Limitations & Workarounds 406

Weighted creases for subdivision surfaces **408**
- Description 408
- Options 408
- Examples 408
- Multiresolution Mesh 409
 - Description 409
 - Options 409
 - Limitations 410
- Workflow 410
- See also 412
- Overview 412
- Sculpt Mode 412
- Hiding and Revealing Mesh 412
- Sculpt Properties Panel 413
- Brushes 413
- Modifiers 414
 - Brush Shape 414
- Brush and Texture Panels 415
- Sculpt Menu 417
- Keyboard Shortcuts 418
- Overview 418
- Usage 418
- Retopo Paint Overview 419
 - Limitations / Notes 421
- Retopo Mesh Overview 421
 - Limitations 421
- Tutorial 421
 - Step One 421
 - Step Two 422
 - Step Three 422
 - Step Four 423
 - Step Five 423
 - Step Six 423
 - Step Seven 424

Curves **425**
- BÃ©ziers 425
 - Curve resolution 427
 - Bevel and Taper Objects 428
- NURBS 428
 - Uniform-Endpoints 429
 - Order 429
 - Weight 430
 - Resolution 431
 - Opening, Closing, Deleting, Joining, Bevel, Taper 431

Table of Contents

Add new segment **432**
Description 432

Opening and Closing a Curve **433**
Description 433
Examples 433

Deleting a Segment(s) **434**
Description 434

Joining two curves **435**
Description 435
Curve Deform 435
Interface 435
Example 437
Curve Taper 439
Important rules to consider 440

Surfaces **442**
Uniform-Endpoints 443
Order 443
Weight 444
Preset Weights 444
Resolution 445
Adding or Extruding 446
Description 446
Examples 446
Cycling (Opening and Closing) 447
Description 447
Deleteing/Erasing surfaces 447
Description 447
Hints 447
Joining or Merging two surfaces 448
Description 448
Examples 448
Hints 448

Skinning **449**
Text 452
Description 452
Options 452
Examples 452
Hints 452
Editing Text 453
Description 453
Options 453
Changing Fonts 454
Description 454
Options 454
Typography 455
Description 455
Options 455

Table of Contents

Skinning
Examples....456
Text Frames....456
Description....456
Options....457
Examples....458
Multiple Materials....459
Description....459
Options....459
Examples....459
Curve and Surface attributes....460
Description....460
Options....460
Examples....460
See Also....461
Special Characters....461
Description....461
Options....461
Unicode Characters....462
Description....462
Technical Details....462
Meta Objects....462
Options....463
Technical Details....463
Threshold (Influence)....464
Description....464
Options....464
â← Positiveâ← & â← Negativeâ← Influence....464
Wiresize....465
Description....465
Options....465
Examples....465
Stiffness....466
Description....466
Options....466
Examples....466
Grouping....467
Description....467
Options....467
Examples....468
Hints....468

Main Section for Scripts....470
Poly Reducer....470
Description....470
Options....470
Hints....471
Examples....471
Auto Image Layout....472
Description....473
Options....473
Examples....473

Table of Contents

Weight Painting **475**
- Bone Weight Copy 475
 - Description 475
- Mirror Vertex Locations & Weight 475
 - Description 475
- Weight Paint Gradient 475
 - Description 476
- Vertex Colour Gradient 476
 - Description 476
 - Examples 476
- Self Shadow 476
 - Description 476
 - Examples 476

Modifiers and deformation **478**
- Interface 478
- Modifiers 478

Modifier **480**

Interface **481**

Stack **483**

Armature Modifier **484**
- Description 484
- Options 484

Cast Modifier **486**
- Description 486
- Options 486
- Examples 487

Curve Modifier **489**
- Description 489
- Options 489
- Example 489

Hooks **490**

Object Hooks **491**
- Description 491
- Examples 491

Adding Hooks **492**
- Description 492
- Options 492

Editing Hooks **493**
- Description 493
- Options 493

Table of Contents

Hook Modifier **494**
Description 494
Options 494
Examples 495
Hook modifier on lattices 495

Lattice Modifier **496**
Options 496
Hints 496
Example/Tutorial(s) 497
Particles and Lattices 498

MeshDeform Modifier **499**
Description 499
Options 499
Other Characteristics of MeshDeform Modifier 501
Mode of Operation 501
Deform Mesh Cage Location after Binding 501
Distance from Deform Mesh Cage & Deformed Object 501
Multires support 502
Interior Control 502
Implementation 503
Examples 503
See Also 503

Shrinkwrap Modifier **504**
Description 504
Options 504
Examples 506

SimpleDeform Modifier **508**
Description 508
Options 508
Notes 510

Wave Modifier **511**
Description 511
Options 511
Technical Details and Hints 513

Array Modifier **514**
Description 514
Options 514
Hints 516
Offset Calculation 516
Examples 516
Mechanical 517
Fractal 517
Organic 518
Tutorials 519

Table of Contents

ParticleInstance Modifier **520**
Description 520
Definitions/Terms 520
Options 521
See Also 526

Build Modifier **527**
Description 527
Options 527
Example 528

Bevel Modifier **529**
Description 529
What is a Bevel? 529
Setting Bevel Weight 530
Options 530
See Also 534

Boolean Tools **535**
Description 535
Options 535

Boolean Modifier **537**
Description 537
Options 538

Examples **539**
Intersect 539
Union 539
Difference 539
More complex examples with Boolean modifiers 540
Multiple Modifiers 540
Animation 540

Technical Details **542**

Limitations & Workarounds **543**

Decimate Modifier **544**
Description 544
Options 544
Examples 545
Simple plane 545
Decimated cylinder 545
High resolution landscape 546

Displace Modifier **547**
Description 547
Options 547
Examples 548
See also 549

Table of Contents

EdgeSplit Modifier **550**
- Description 550
- Options 550
- Examples 551

Explode Modifier **552**
- Description 552
- Options 552
- Examples 557
 - Exploding Cube 557
 - Another explosion 557

Mask Modifier **558**
- Description 558
- Options 558
- Examples 560

Mirror Modifier **561**
- Description 561
- Options 561
- Hints 563
 - Using Mirror modifier with Subsurf modifier 564
 - Aligning for Mirror 564

Smooth Modifier **565**
- Description 565
- Options 565
- Examples 565

Subdivision Surfaces **567**
- Description 567
- Options 567
- Hints 568
- Examples 568
- Limitations & Workarounds 572

Weighted creases for subdivision surfaces **574**
- Description 574
- Options 574
- Examples 574

UVProject Modifier **576**
- Description 576
- Options 576
- Examples 577

Introduction **579**
- Viewing Restrictions 579
- Global Influences 579
- Lighting Settings 579
- Lighting in the Workflow 580
 - Over-riding Materials to reset lighting 580

Table of Contents

Lights **581**

Lights Common Options **582**

Light Attenuation **583**
- Description 583
- Falloff types 583
 - Lin/Quad Weighted 583
 - Custom Curve 586
 - Inverse Square 588
 - Inverse Linear 588
 - Constant 589
- Sphere 589
- Examples 590
 - Distance 590
 - Inverse Square 591
 - Sphere 592
- Hints 593
- See Also 593

Lamps Textures **594**
- Options 594

What the Light Affects **595**

Lamps Related Settings **596**
- Lighting Groups 596
 - Materials 596
 - Scene 596
- Other Options 597
- See Also 597

Shadows **598**
- Lamps: Raytraced Shadows 598
- Lamps: Buffered Shadows 599

Common Shadowing Lamps Options **601**
- See Also 602

Lamps Raytraced Shadows **603**

What is Quasi-Monte Carlo? **606**

Volumetric Lighting **607**

Lamps **608**

Lamp **609**
- Description 609
- Options 609
- â← Shadow and Spotâ← Panel 610
- Examples 610

Table of Contents

Lamp
See Also....610

Lamp Raytraced Shadows....611
See Also....611

Spot....612
Description....612
Options....612
â← Shadow and Spotâ← Panel....613
Options....613
Shadows and Halos....615
Technical Details....615
Examples....616
See Also....616

Spot Raytraced Shadows....617
See Also....617

Spot Buffered Shadows....618
Buffered Shadows....618
â← Classicalâ← and â← Classic-Halfwayâ← generation methods....619
â← Irregularâ← generation method....625
Common options....626
Note....629
Examples....629
Hints....630
See Also....630

Spot Volumetric Effects....631
See Also....633

Area....634
Description....634
Options....634
â← Shadow and Spotâ← Panel....635
Examples....636
Technical Details....636
Hints....637
See Also....637

Area Raytraced Shadows....638
Raytraced Shadows....638
See Also....639

Hemi....640
Description....640
Options....640
â← Shadow and Spotâ← Panel....641
Examples....641
Outdoor Light....641
Technical Details....642

Table of Contents

Hemi
See Also 642

Sun **643**
Description 643
Options 643
â← Sky/Atmosphereâ← Panel 644
â← Shadow and Spotâ← Panel 644
Examples 644
Hints 645
See Also 645

Sun Raytraced Shadows **646**
See Also 646

Sun: Sky & Atmosphere **647**
Description 647
Options 648
Sky 648
Atmosphere 649
Examples 649
Sky 650
Atmosphere 653
Hints and limitations 655
See Also 655

Lighting Rigs **656**
Ambient Only 656
Single Rig 657
Two-Point Rig 658
Three-Point Rigs 658
Studio rig 659
Standard Rig 659
Four-point Rig 660
Troubleshooting and a Helpful Blend file 661

Introduction **662**
Technical Details 662

Radiosity Rendering **665**
Description 665
Enabling Radiosity 665
Material Options 666
Rendering Options 666
Examples 667
Hints 668

Baking Radiosity **669**
Radiosity Tool 669
Description 669
Collecting Meshes 669
Description 669

Table of Contents

Baking Radiosity
Options....669
Subdivision limits....670
Description....670
Options....671
Adaptive Subdivision....671
Description....671
Options....671
Computing Radiosity....672
Description....673
Options....673
Editing the solution....674
Description....674
Options....674

Ambient Light....676
Description....676
Setting Ambient Light....676
Material Ambient Effect....677
World Ambient Occlusion....677
Radiosity....677

Ambient Occlusion....678
Description....678
Options....678
Raytrace....679
Approximate....681
Common Settings....682
Technical Details....683
Hints....683

Exposure and Range....684
Description....684
Options....684
Examples....684
Hints....685

Introduction to Shading....686
Basics: Materials and Textures....686
Advanced: Vertex painting and UV texturing....686

Introduction....688
Creating a new Material....689
Sharing a Material from another object....689
Materials Options....689

Material Preview....691
Description....691
Options....691
Examples....691

Table of Contents

Materials **693**
Description 693
Links and Pipeline Options 693
The Material Index 695
Assigning Material Indexes to a Mesh 696
Adding a Material Index 697
Deleting a Material Index 697
Assigning Face(s) to a Material Index 697
Changing a Material used by a Material Index 698
Checking your work using Select, Deselect, Query (?) 698
Some Examples 699
Advanced Multiple Materials 699

Diffuse Shaders **700**
Description 700
Options 700
Technical Details 700
Hints 701

Lambert **702**
Description 702
Options 702

Oren-Nayar **703**
Description 703
Options 703

Toon **704**
Description 704
Options 704

Minnaert **706**
Description 706
Options 706

Fresnel **707**
Description 707
Options 707

Specular Shaders **709**
Description 709
Options 709
Technical Details 709

CookTorr **711**
Description 711
Options 711

Phong **712**
Description 712
Options 712
Planet Atmosphere 712

Table of Contents

Blinn **713**
Description 713
Options 713

Toon **714**
Description 714
Options 714
Hints 714

WardIso **715**
Description 715
Options 715
Material Ambient Effect 715
Ramps 716
Description 716
Options 716
Colorbands 718
Description 718
Options 719
Examples 720

Raytraced Mirror Reflections **723**
Description 723
Options 723
Colored reflections 725
Examples 726
Reference 727
Hints 727

Raytraced Transparency **728**
Description 728
Options 728
Examples 729
Hints 733
How it works 737
Enabling SubSurface Scattering 738
Options 739
Developing your own SSS material 740
Example: Grapes 740
Example: Skin 741

Strands **743**
Strands Shading 743
Texturing along the Strand 745
Strand render Simplification 746
Introduction 746
Accessing The Node Editor 747
Enabling Node Materials in the Material Buttons 748
External Links 748
Accessing The Node Editor 749
Activating Nodes 750
Required settings for Composition 750

Table of Contents

Strands
Node Editor Window Actions....751
Node Editor Header....751
View, Select, and Add Header Menus....752
Node Header Menu....752
Material/Composting Selector Button....753
Use Nodes Header Button....753
Free Unused Header Button....753
Node Basics....753
What are nodes....753
Adding and Arranging Nodes....754
Sockets....754
Connecting and Disconnecting Sockets....755
Node Groups....755
Node Controls....757
Node Sizing....758
Node Curves....759
RGB Curves....759
Selecting curve points....759
Editing curves....760
Editing the view....760
Special tools....760

Material Input Nodes....761
Material Node....761
Output....761
Input....762
Controls....762
Using the Material Node with Specularity....762
Value Node....763
RGB Node....763
Texture Node....764
Geometry Node....764
Geometry Node Example using a UV image....765

Material Output Nodes....766

Material Color Nodes....767
Mix Node....767
Using Dodge and Burn (History Lesson)....769
RBG Curves Node....769

Material Vector Nodes....771
Mapping Node....771
Mapping Node Example....771
Normal Node....771
Normal Node Example....772
Using Normal to reflect bumpy-looking textures....773
Using the Normal Node to create special Light-Sources....774

Table of Contents

Material Convertor Nodes **776**
ColorRamp Node 776
RGB to BW Node 776

Groups of Nodes **778**
Create a Node Group 778
Editing Node Groups 778
Ungroup Nodes 778

Dynamic Nodes (PyNodes) **779**
Setting Up 779
Color Picker Applet 780
Vertex Paint Panel 781
Saving and Exporting 781
Overpainting and Textures 782

Halo Materials **783**
Options 783
Halo Texturing 784
Examples 785
Dotmatrix display 785
Lens Flares 786

Introduction **788**

Texture Buttons **790**
Description 790

Texture Preview **791**
Description 791
Options 791

Texture Channels **792**
Description 792
Options 792

Texture Colors **793**
Description 793
Options 793
Hints 794
See Also 794

Texture Channels **795**
Description 795
Channels Options 795
The channels stack 796
(De)activating a channel 797
Changing the order of the textures 797
Copying and pasting Texture Channels 797
Texture Map Input 797
Description 797
Input Source 798

Table of Contents

Texture Channels
Description....798
Options....798
Hints....800
2D to 3D Mapping....800
Description....800
Options....801
Coordinate Offset, Scaling and Transformation....802
Description....803
Options....803
3D View texture space transform....804
Description....804
Options....804
Examples....805
Position a Decal on a Mesh....805

Texture Map To....806
Description....806
Options....806
Examples....808
Hints....809

Stencil....810
Description....810
Options....810
Examples....810

Warp....812
Description....812
Options....812
Examples....812

Bump and Normal Maps....814
Description....814
Workflow....814
Examples....815
Baking Options....816
Using And Creating Normal Maps....817

Displacement Maps....819
Description....819
Options....819
Hints....819
Using the Displace Modifier....820
Examples....820
See Also....822
How to create a Displacement Map....822

Procedural Textures....823
Description....823
Options....823
Hints....823

Table of Contents

Noise Basis **825**
Description 825
Examples 825

Clouds **826**
Description 826
Options 826
Examples 826
Technical Details 827

Marble **828**
Description 828
Options 828
Examples 828
Technical Details 829

Stucci **830**
Description 830
Options 830
Examples 830
Technical Details 831

Wood **832**
Description 832
Options 832
Examples 832
Technical Details 833

Magic **834**
Description 834
Options 834
Examples 834
Technical Details 835

Blend **836**
Description 836
Options 836
Examples 836
Technical Details 837

Noise **838**
Description 838
Options 838
Technical Details 838

Musgrave **839**
Description 839
Options 839
Examples 840
Technical Details 840

Table of Contents

Voronoi **841**
Description 841
Options 841
Examples 842
Technical Details 842

Distorted Noise **843**
Description 843
Options 843
Examples 843
Technical Details 843

Image Textures **844**
Description 844
Map Image Panel Options 844
Image Panel Options 847
Seamless Images 848

Video Textures **850**
Description 850
Options 850
Examples 851
Limitations 852

Environment Maps **853**
Description 853
Options 853
Examples 855
Limitations 856
Introduction 856
Using Texture Nodes 857
Using Multiple Outputs 857
Setup examples 858
See also 859

Texture Plugins **860**
Description 860
Sample List of Available Texture Plugins 860
Options 862
Technical Details 862
See Also 862
UV Explained 862
Cartography Example 863
Half-Sphere Example 864
Advantages 864
Getting Started 865
Workflow 865
UV Orientation 867
UV Rotation 867
UV Location 868
Mapping an Image to Your UVs 868
Images versus Maps 870

Table of Contents

Texture Plugins
Panel Interaction....870
Unwrapping Using Seams....871
Unwrapping Multiple Faces....873
Face Unwrap using UV Calculation....875
Activating UV Textures....879
Combining UV Maps....880
Averaging UV Islands....882
Multiple UV Layouts....882
Additional Options....883
Editing Buttons....883
Saving your UV Layout....885
Saving an outline of your UV Layout....885
Major Adjustments....886
3D View: Face Mirror and Rotate UVs....886
Merging UV Maps....886
Separating UV Maps....887
Arranging the UV Layout....887
Deleting a UV Layout....887
UV/Image Editor Menu....888
View Menu....888
Select Menu....889
Image Menu....890
UVs Menu....890
Iteration and Refinement....891
Tradeoffs....892
Using the Pin command....892
Refining the Layout....892
Reusing Textures....893
Goals....893
Workflow....894
Using Images and Materials....894
Creating an Image Texture....894
Creating a New Blank Image....895
Modifying your Image Texture....896
Replacing an image via Render Bake....897
Creating an image via Texture Baker Script....897
Create/Modify via an External Image Paint Program....898
Modify an image via Texture Paint....898
Using a Saved Image....898
Mapping the Image Texture....899
Replacing the active Image....900
Packing Images inside the Blend file....900
Layering UV Textures....900
Mix and Match Materials....901
Using Alpha Transparency....902
UV Textures vs. Procedural Textures....903
UV/Image Editor Menu....904
Window Header Buttons....904
Image Menu....905
UVs Menu (Image Related)....906

Table of Contents

Using Blender's Texture Paint **907**
- Getting Started 907
 - Instant Feedback 908
- Brushes 909
 - Brush Mode Options 909
 - Mix Mode Options 910
 - Brush Control 911
 - Brush Pattern/Texture 911
- Saving 912
- Options 912

Using an External Image Editor **913**

2.49 Release **914**
- Projection Painting 914
 - Features 914
 - Implementation 917
 - Limitations 917
 - Future Plans 917

World options **919**

World Background **920**
- Description 920
 - Options 920
 - Textures 921

Mist **922**
- Description 922
 - Option 922
 - Examples 923

Stars **924**
- Description 924
 - Options 924
 - Examples 924
- Chapters 926
- Interpolation in Blender 926
- Ipo Curve Types 927
- Animating Materials 928
 - Creating Channels 928
 - Animating Channel Values 929
 - Animating Texture Channel Mapping 930
 - Driving Channel Values 930
- Animating Lamps 930
- Animating World Settings 931
- Animating Constraints 931
- Sharing Ipo Datablocks 931
- See Also 932
 - Example 932
- Recording IPOs 933
- Linking IPOs 933

Table of Contents

Ipo Curves **934**
- Curve showing/hiding 934
- Curve Selection 935
 - Options 935
- Curve Manipulation 935
 - Options 935
- Curve Interpolation 935
 - Options 936
- Curve Extend Mode 936
 - Options 936
 - Examples 936
- Editing Curve Points 937
- Adding Curve Points 937
 - Hints 937
- Point Handle Types 938
 - Options 938

Ipo Keys **939**
- Onion Skinning and Ghosting 939
 - Hints 939
 - Examples 940

Special Notes on the Time Curve **941**
- Examples 941

Timeline Window **942**
- View Menu 942
- Frame Menu 942
- Playback Menu 942

Markers **944**
- Creating Markers 944
- Naming Markers 944
- Moving Markers 944
- Deleting Markers 945
- Creating a Path 945
 - Path Buttons Panels 946
 - Changing the Path 947
- Following a Path 947
 - Speed/Progress via Time Curve 947

Legacy Path Method **949**
- Tips 952
- Description 953
- Workflow 953
- Creating Relative Shape Keys 954
- Editing Shape Keys 955
- Pinning 956
 - Hints 956
- Blending Shape Keys 956
 - Blending Using the Shapes Panel 957
 - Blending Using the Action Editor 957

Table of Contents

Legacy Path Method
Blending Using the Ipo Curve Editor....957
Using IPO Drivers....958

Absolute Shape Keys....959
Description....959
What Types of Objects can be Shaped?....959
Creating Absolute Shape Keys....959
Sharing Shape Keys....961
Workflow for Creating Absolute Shape Keys....961
Shifting between Shapes....962
Options....963
Hints....963
Curve and Surface Keys....964
Lattice Animation....964
Implications for Camera Tracking....965
IPO Header....966
WindowType....966
Menus....966
Extend mode Constant....966
Extend mode Direction....966
Extend mode Cyclic....967
Extend mode Cyclic Extrapolation....967
IPO Type....967
IPO Menu....971
Pin IPO....972
IPO Menu....972
IP :....972
Users....972
Unlink IPO....972
Fake User....972
Copy to Buffer....972
Paste from Buffer....972
View Border....973
Lock....973
IPO Window....973
IPO Curve Select....974
Channel Select....974
Mouse....974
CtrlLMB....974
MMB....974
CtrlMMB....974
RMB....974
RMB and drag....975
â § ShiftRMB....975
HotKeys....975
- NumPad + NumPad....975
Pageup....975
â § ShiftPageup....975
Pagedown....975
â § ShiftPagedown....975
â Home....975

Table of Contents

Absolute Shape Keys
â Tab....976
A....976
B....976
â § ShiftC....976
C....976
â § ShiftD....976
G....976
H....977
â § ShiftH....977
I....977
J....977
K....977
R....977
S....978
â § ShiftS....978
T....978
V....979
X....979
See also....979

Hooks....980

Object Hooks....981
Description....981
Examples....981

Adding Hooks....982
Description....982
Options....982

Editing Hooks....983
Description....983
Options....983

Hook Modifier....984
Description....984
Options....984
Examples....985
Hook modifier on lattices....985

Introduction To Blender Constraints....986
General Constraint Properties....986
Influence....986
Multiple Influence Example....987
Constraints on Bones....987

Constraint Types....988
Copy Location....988
Copy Rotation....988
Copy Scale....989
Limit Constraints....990

Table of Contents

Constraint Types
Track To....991
Settings....991
Floor....991
Settings....992
Locked Track....992
Settings....992
Follow Path....993
Settings....994
Stretch To....994
Settings....995
IK Solver....995
Settings....995
Action....996
Null....996

Character Animation using the Armature System....997
Chapters....997

Adding an Armature....999
Description....999
Options....999

Armature Datablock....1000
Description....1000
Options....1000

Armature Display Options....1001
Description....1001
Options....1001

Armature Deformation....1004
Description....1004
Options....1004

Editing Armatures....1005
Description....1005
Options....1005

Selecting Bones....1008
Description....1008

Adding and Deleting Bones....1009
Description....1009

Interactively Setting Bone Roll....1010
Description....1010

Subdivide....1011
Description....1011

Table of Contents

X-Axis Mirror Edit **1012**
Description 1012

Flipping Left and Right Bone Names **1013**
Description 1013

Naming conventions **1014**

Armature Posing Options **1016**
Description 1016
Options 1016

Armature Posing Tools **1018**
Description 1018

Inverse Kinematics **1019**
Interface 1019
Tree IK 1019
IK target rotation 1019
IK Pole Target 1020
Rotation limits 1020
Setting rotation range 1020
Mixing IK and FK 1021

Skinning **1022**
Weight Painting 1023
Context menu 1024

The Action Editor **1025**
Long Keyframes 1025
Shape Key Sliders 1025
Action Window 1025
Selection Tools 1025
Mirror Tools 1026
Snap Tools 1026
Shift-S Snapping 1026
Auto-Snapping During Transforms 1027
Channel 'Protecting' 1027
Tip 1027

NLA Editor **1030**
Description 1030

NLA Strip Options **1032**
Description 1032
Properties 1033

NLA Strip Tools **1035**
Description 1035
Options 1035
Interface 1036
Constraints 1036

Table of Contents

NLA Strip Tools
Influence....1037
Multiple Influence Example....1037
Constraints on Bones....1038

Child Of Constraint....1039
Description....1039
Options....1039
Tips....1039
Example....1040

Copy Location Constraint....1041
Description....1041
Options....1041
Example....1041

Copy Rotation Constraint....1043
Description....1043
Options....1043
Example....1043

Copy Scale Constraint....1045
Description....1045
Options....1045
Example....1045

Limit Location Constraint....1046
Description....1046
Options....1046
Example....1047

Limit Rotation Constraint....1048
Description....1048
Options....1048
Example....1049

Limit Scale Constraint....1050
Description....1050
Options....1050
Example....1050

Floor Constraint....1051
Description....1051
Options....1051
Example....1051

Track To Constraint....1052
Description....1052
Options....1052
Example....1052
Piston....1053

Table of Contents

Locked Track Constraint **1054**
Description 1054
Options 1054
Example 1055

Follow Path Constraint **1056**
Description 1056
Options 1056
Example 1057

Clamp To Constraint **1058**
Description 1058
Options 1058
Example 1059

Stretch To Constraint **1060**
Description 1060
Options 1060
Example 1060

Transformation Constraint **1061**
Description 1061
Options 1061
Example 1062

Rigid Body Joint Constraint **1063**
Description 1063
Options 1063
Example 1063

IK Solver Constraint **1064**
Description 1064
Options 1064
Example 1065

Action Constraint **1066**
Description 1066
Options 1066
Example 1066

Script Constraint **1067**
Description 1067
Options 1067
Example 1067

Introduction **1069**

Driven Shape Keys **1071**
Setup the Shape Keys 1071
Connecting Shape key and Driver 1073
Custom Controls 1074
NLA 1076

Table of Contents

Driven Shape Keys
Define Actions....1076
Links....1077

Introduction....1078
Description....1078
Workflow....1079
Creating a Particle System....1080
Types of Particle systems....1080
Common Options....1081
Bake....1081
Interaction in real time....1082
Links....1082
Emitter....1083
Options....1084
Usage....1084
Reactor....1085
Options....1085
Usage....1086
Hair....1087
Options....1087
Usage....1088
Physics: None....1089
Physics: Newtonian....1089
Integrators....1089
Initial velocity....1089
Rotation....1090
Global effects....1090
Physics: Keyed....1091
Physics: Boids....1092
Behaviour....1092
Physics....1093
Boids, deflectors and effectors....1094
Links....1094
Types of Visualization....1095
Draw....1097
Render....1097
Path....1097
Billboard....1098
Time....1100
The Extras Panel....1101
Effectors....1101
Time....1101
Vertex group....1102
Other....1102
Textures controlling particle attributes....1103
Ipo channels....1104
Controlling particle parameters....1104
Particles....1105
Kink/Branch....1107
Faces....1107
How to get the Particle Mode up and running....1108

Table of Contents

Introduction
Particle Mode Options....1108
Selecting keypoints....1108
Moving keypoints or particles....1109
Mirroring particles....1109
Hiding/Unhiding....1109
Select Modes....1109
The Particle Edit Properties Panel....1109
Particle system....1111
Add volume and dress....1112
Material....1113
Black hair....1113
Blond hair....1114
Animation....1114
Links....1114
Typical scenarios for using Soft Bodies....1115
Creating a Soft Body....1116
Cache and Bake....1116
Interaction in real time....1117
Tips....1117
Links....1118
Exterior Forces....1118
Keeping the vertices at their place (â← Goalâ←)....1119
Edges....1121
Stopping an object from collapsing....1122
Attribution....1123
Collisions with other objects....1124
Motivating examples....1124
Calculating the collisions....1124
Good collisions....1125
Self Collision....1126
Balls....1126
Flag in the wind....1127
Interaction....1129
Swing....1131
Goal, Goal, Goal!....1133
Turning jello into hair....1134
Collisions and Deflection....1135
Soft Body settings....1136
General settings....1136
Goal settings....1137
Edge settings....1137
Bake and Collision....1138
Bake....1138
Self Collision....1139
Soft Body Solver....1140
Links....1140
Spherical field....1141
Other Field Types....1143
Texture field....1144
Curve Guide field....1145
Links....1147

Table of Contents

Introduction
Options....1147
Particle Interaction....1148
Soft Body and Cloth Interaction....1148
Effector Interaction....1149
Examples....1149
Hints....1149
Setting up world physics....1149
Setting up your scene....1150
Setting up collision objects....1150
Animation preview....1151
Get ready to bake....1152
Record the Ipo....1152
Scrub through the animation....1152
Replace the animation....1153
Troubleshooting and hints....1153

Fluid Simulation....1154
Description....1154
Process....1154
Options....1155
Domain....1155
Fluid....1160
Obstacle....1161
Inflow....1162
Outflow....1163
Particle....1163
Control....1164
Animating Fluid Property Changes....1164
Technical Details....1165
Hints....1167
Limitations & Workarounds....1168
See Also....1169
Acknowledgements....1169

Cloth Simulation....1170
Description....1170
Workflow....1171
Cloth Panel Options....1171
Cloth Definition....1171
Using Simulation to Shape/Sculpt a Mesh....1172
Collisions....1172
Bake Collision....1175
Editing the cached simulation....1175
Troubleshooting....1176
Examples....1176
Pinning of cloth....1177
Smoothing of Cloth....1177
Cloth on armature....1177
Using Cloth for Softbodies....1177
Cloth with Wind....1178

Table of Contents

Introduction **1179**
Overview 1180
Render Workbench Integration 1180
Sequencer from/to Compositor 1181
Render Window Options 1181
Showing Previous Renders 1182
Render Window usage 1182
Step Render Frame 1183
Distributed Render Farm 1183
Show the Render Window 1184
Rendering 1184
Showing 1184
Render Window usage 1185
See also 1185

Shadow **1186**

EnvMap **1187**

Raytracing **1188**
Octree resolution 1188

Radiosity **1189**

Percentage Size **1190**

Parts Rendering **1191**
Description 1191
Options 1191
Limitations & Work-arounds 1192

Panoramic Rendering **1193**
Description 1193
Options 1193
Examples 1194

Motion Blur Rendering **1196**
Description 1196
Options 1196
Examples 1196
Hints 1197

Stamp **1198**

Oversampling **1199**
Description 1199
Options 1199
Examples 1201

Rendering with Yafray **1204**
Description 1204
Options 1204

Table of Contents

Rendering with Yafray
Examples....1206

Notes....1209
Amount of Light....1209
SkyDome....1210

Render Baking....1211
Description....1211
Options....1212
Options under development (after 2.45)....1212
Workflow....1213
Refining Baked Images....1214
Examples....1214
Applying Baked Images....1216
Refining the Ambient Occlusion Image....1217
Applying High-Res Normal to a Low-Res Mesh....1218
Hints....1218
Blender to xNormal/Maya Normal....1219

Introduction....1220

Description....1221

Examples....1223
Render a picture....1223
Render a movie....1223

Output Panel....1224
File Locations....1224
Scene Linking....1225
Other Buttons and Options....1225

Edge (Toon) Rendering....1227
Description....1227
Options....1227
Examples....1228

Output Formats....1229
Description....1229
Format Panel....1229
Presets....1229
Options....1230
FFMPEG....1233
Free Advice....1234

Checking Render Results....1236

Rendering Animations....1237
Description....1237
Workflow....1237
Options....1239

Table of Contents

Rendering Animations
Hints....1240

Preparing your work for video....1241
Colour Saturation....1241

Rendering to fields....1243
Description....1243
Options....1243

Finding Render Layers....1245

Using Render Layers....1246

Render Layer settings....1247
Enabling and Naming....1247
Scene Layers to Render....1247
Pipeline Products to deliver....1247
Passes to make available....1248

Creating a new Render Layer....1249

Examples....1250
Rendering only certain objects....1250
Outlining only selected objects....1250
RenderPass In Detail....1250
Selecting RenderPasses....1251
De-Selecting Render Passes....1253
Using RenderPasses....1253
Recoloring Shadows....1253
Compositing Ambient Occlusion....1254
Vector Blurring Shadows....1255
Conclusion....1255

Vanishing points in Render....1257
Three Point Rendering....1257
Two Point Rendering....1258
One Point Rendering....1259
Zero Point (Orthographic) Rendering....1260
Isometric Rendering....1260

Depth Of Field (DOF) Explained....1262
Field of View and Lens Size....1262
Zooming in Blender....1263
Depth of Field in Computer Graphics....1263
Rendering Big Hairy Scenes or Speeding Up Rendering....1263
Hardware Improvements....1264
Operating System Configuration....1265
Blender Choices....1265
Scene and Specific Objects....1265
Render Settings....1266
Another Way To Do It: Compositing....1267

Table of Contents

Introduction **1268**

Accessing and Activating Nodes **1270**

Grouping Compositing Nodes **1271**
- Create a Node Group 1271
 - Editing Node Groups 1271
 - Ungroup Nodes 1271

Some Examples of Using Compositing Nodes **1272**
- Accessing The Node Editor 1272
- Activating Nodes 1273
 - Required settings for Composition 1274
- Node Editor Window Actions 1274
- Node Editor Header 1275
 - View, Select, and Add Header Menus 1275
 - Node Header Menu 1276
 - Material/Composting Selector Button 1276
 - Use Nodes Header Button 1277
 - Free Unused Header Button 1277
- Node Basics 1277
 - What are nodes 1277
 - Adding and Arranging Nodes 1277
 - Sockets 1278
 - Connecting and Disconnecting Sockets 1279
 - Node Groups 1279
- Node Controls 1280
- Node Sizing 1282
- Node Curves 1282
 - RGB Curves 1283
 - Selecting curve points 1283
 - Editing curves 1283
 - Editing the view 1283
 - Special tools 1284
- Render Layers Node 1284
 - Using the Alpha Socket 1285
 - Optional Sockets 1286
 - Using the Z value Socket 1287
 - Using the Speed Socket 1288
- Image node 1288
 - Image Channels 1289
 - Saving/Retrieving Render Passes 1289
 - Image Size 1290
 - Animations 1290
 - Generated Images 1291
 - Notes 1292
 - Splicing Video Sequences using Nodes 1292
- Texture Node 1295
 - Example 1296
- Value node 1296
- RGB node 1297
 - Example 1297

Table of Contents

Some Examples of Using Compositing Nodes
Time node....1297
Time Node Examples....1298

Using Input Nodes....1299
Viewer....1299
Using the UV/Image Editor Window....1299
Composite....1300
Saving your Composite Image....1300
SplitViewer Node....1300
File Output Node....1301
RGB Curves Node....1302
Options....1302
Examples....1303
Mix Node....1304
Examples....1306
Contrast Enhancement using Mix....1307
Using Mix to Watermark images....1308
Using Dodge and Burn (History Lesson)....1309
Hue Saturation Node....1309
Hue Saturation tips....1310
HSV Example....1310
Bright/Contrast....1311
Notes....1311
Gamma....1312
Invert....1312
Options....1313
AlphaOver Node....1313
Examples....1314
Strange Halos or Outlines....1314
Z-Combine Node....1315
Examples....1316
Normal Node....1317
Vector Curves Node....1317
Map Value Node....1318
Using Map Value to Multiply values....1318
Filter Node....1319
Blur Node....1320
Options....1320
Example....1320
Bilateral Blur Node....1321
Inputs....1321
Options....1321
Examples....1321
Vector (Motion) Blur Node....1322
Examples....1323
Known Bugs....1323
Dilate/Erode Node....1323
Example....1324
Defocus....1324
Camera Settings....1324
Node Inputs....1325

Table of Contents

Using Input Nodes
Node Setting....1325
Examples....1327
Hints....1328
ColorRamp Node....1329
Using ColorRamp to create an Alpha Mask....1329
Using ColorRamp to Colorize an Image....1330
RGB to BW Node....1330
Set Alpha Node....1331
Using SetAlpha to Fade to Black....1331
Using SetAlpha to Fade In a Title....1332
Using SetAlpha to Colorize a BW Image....1332
ID Mask Node....1333
Example....1334
Math Node....1334
Examples....1335
Combine/Separate Nodes....1339
Separate/Combine RGBA Node....1339
Separate/Combine HSVA Nodes....1340
Separate/Combine YUVA Node....1341
Separate/Combine YCbCrA Node....1341
Difference Key Node....1343
Simple Example....1344
Chaining Example....1344
Chroma Key Node....1345
Luminance Key Node....1346
Example....1346
Color Spill Node....1347
Channel Key Node....1347
Translate Node....1348
Example: Using the Translate Node to Roll Credits....1349
Example: Moving a Matte....1350
Example: Shake, Rattle and Roll....1350
Rotate Node....1352
Scale Node....1352
Flip Node....1352
Displace Node....1353
Example....1353
Map UV Node....1354
Examples....1355

Introduction....1356
Chroma Vectorscope....1357
Luma Waveform....1357
Image Preview....1358
Sequence Mode....1359
Overview of the Video Sequence Editor Window....1359
View Menu....1360
Select Menu....1361
Add Menu....1362
Add Movie and Image (Stills or Sequences)....1362
Add Scene....1364

Table of Contents

Introduction
Add Audio....1364
Add Effect....1364
Animating Audio/Video Effects....1366
Strip Menu....1367
Strip Properties....1369
Input Strip Properties Panel....1371
Filter Strip Properties Panel....1371
Proxy Strip Properties Panel....1371
Edit Strip Properties Panel....1371
Working with Strips....1372
Scrubbing....1372
Selecting and Changing the Length of a Strip....1373
Zoom, Scale Display and Refresh Header Buttons....1373
Reverse Action....1374
Slow Motion, Fast Forward, Time Warp....1374
Rendering a Video to an Image Sequence Set....1374
Rendering a Video from an Image Sequence Set....1374
Add Effect....1375
Subtract Effect....1375
Multiply Built-in Effect....1376
Cross and Gamma Cross....1376
Alpha Over, Under, and Over Drop Built-in Effects....1377
Wipe Built-in Effect....1379
Glow Built-in Effect....1379
Transform Built-in Effect....1380
Color Generator Built-in Effect....1380
Speed Control Built-in Effect....1381
Alpha Matte Plugin Effect....1383
Animated Glow Plugin Effect....1383
Box Blur Plugin Effect (Pre 2.48 only; crashes 2.48+)....1384
Cartoon Plugin Effect....1384
Chromakey RGB Plugin Effect....1384
Chromakey YUV Plugin Effect....1385
Clockwipe Plugin Effect....1385
Desaturate Plugin Effect....1386
Dice Plugin Effect....1386
Diff Plugin Effect....1387
Displace Plugin Effect....1387
FlipFlop and FlipIt Plugin Effects....1388
GreenScreen Plugin Effect....1388
Iris Plugin Effect....1389
Jitter Plugin Effect....1390
Lightsaber Plugin Effect....1390
LSD Plugin Effect....1391
OldMovie Plugin Effect....1391
Red-Green Stereo Plugin Effect....1392
Scatter Plugin Effect....1392
Warp Plugin Effect....1393

Table of Contents

Sound Sequence Editor **1394**
Options 1394
Audio Mixing in the VSE 1394
Working with Audio Tracks 1395
Animating Audio Track Properties 1395

Extending Blender **1396**
Python scripts 1396
Plugins 1396
Saving your own Extensions 1396
Files location 1396
Scripts 1397
Plugins 1397

Python Scripts **1398**

Python Scripting **1399**

Downloading Python **1400**

Setting the PYTHONPATH environment variable **1401**
Win95,98,Me 1401
WinNT,2000,XP 1401
Linux and other UNIXes 1401

The Text Editor **1403**
Other usages for the Text window 1404

A working Python example **1405**
Headers, importing modules and globals 1405
Drawing the GUI 1406
Managing Events. 1406
Mesh handling 1407
Conclusions 1408

Blenderâ ← s Python API **1410**

Scripts **1411**

Bundled Scripts **1412**
User Preferences Window 1412
3D View Window 1413
UV/Image Editor Window 1414
Text Editor Window 1414
Making a Script â ← Bundledâ ← 1414
Additional Resources 1415

Blenderâ ← s Plugins System **1416**

Writing a Texture Plugin **1417**
Specification 1417
Texture/Material Interaction 1419

Table of Contents

Writing a Texture Plugin
Generic Texture Plugin....1420
Our Modifications....1421
Compiling....1422

Writing a Sequence Plugin....1424
Specification....1424
Generic Sequence Plugin....1426
Our Modifications....1427
Compiling....1428

Blender Game Engine Overview....1429
Using the Game Engine....1429

Game Logic....1430
1 Associated Object(s)....1430
2 Physics....1431
3 Properties....1431
4 Sensors....1431
5 Controllers....1431
6 Actuators....1432
7 Links....1432

Object types....1433

Object physics in case of no active World....1434

Occluder object type....1435
Invisible....1435

Technical specifications....1436
View frustrum culling....1436
Occlusion culling....1436
How it works....1436

No collision object type....1439
Invisible....1439

Sensor object type....1440
Detect actor....1440
Invisible....1440
Advanced....1440
Bounds....1440

Technical specifications....1441
Sensor Object Type....1441

Static object type....1443
Actor....1443
Ghost....1443
Invisible....1443
Advanced....1443

Table of Contents

Static object type
Anisotropic....1443
Bounds....1443

Dynamic object type....1445
Actor....1445
Ghost....1445
Invisible....1445
Advanced....1445
Mass....1445
Radius....1446
No Sleeping....1446
Damp....1446
RotDamp....1446
Do Fh....1446
Rot Fh....1446
Form....1446
Anisotropic....1447
Bounds....1447

Rigid body object type....1448
Actor....1448
Ghost....1448
Invisible....1448
Advanced....1448
Mass....1448
Radius....1449
No Sleeping....1449
Damp....1449
RotDamp....1449
Do Fh....1449
Rot Fh....1449
Form....1449
Anisotropic....1450
Bounds....1450

Soft body object type....1451
Actor....1451
Ghost....1451
Invisible....1451
Advanced....1451
Mass....1451
Shape match....1451
Bending Const....1452
LinStiff....1452
Friction....1452
kMT....1452
Bounds....1452

Properties....1453
Properties Panel....1453
Property types....1453

Table of Contents

Properties
Property Sensor....1454
Property Actuator....1454

Sensors....1456
Sensor viewing options....1456
General object options....1456
General sensor options....1457
Sensor header....1457
Sensor pulses bar....1457

Sensor types....1460
Always sensor....1460
Delay sensor....1460
Keyboard sensor....1460
Mouse sensor....1461
Touch sensor....1462
Collision sensor....1462
Near sensor....1462
Radar sensor....1463
Property sensor....1463
Random sensor....1464
Ray sensor....1464
Message sensor....1465
Joystick sensor....1465
Actuator sensor....1465

Controllers....1466
Controller Comparison table....1466
AND Controller....1466
Example....1467
OR Controller....1467
Example....1467
XOR Controller....1467
Example....1467
NAND Controller....1467
Example....1468
NOR Controller....1468
Example....1468
XNOR Controller....1468
Example....1468
Expression Controller....1468
Python Controller....1468
Script Mode....1469
Module Mode....1469

Actuators....1470

Motion Actuator....1471
Simple Motion....1471
Servo Control....1472

Table of Contents

Shape Action Actuator **1473**
Options 1473

Action Actuator **1474**
Options 1474

Ipo Actuator **1475**
Example 1476

Property Actuator **1477**
Example 1477

Edit Object Actuator **1478**

Scene Actuator **1479**

Actuator: 2D Filters **1480**
Motion Blur 1480
Built-In 2D Filters 1481
Custom Filters 1482

Camera in BGE **1483**
Default camera 1483
Parent Camera to Object 1483
Parent to Vertex 1483
Object as a Camera 1483
See also 1483

Dome Mode in Blender Game Engine **1485**
Remember to use Blender in fullscreen mode to get the maximum out of your projector. 1485

GUI options: **1486**

Dome Modes **1487**
1. Full Dome Mode 1487
2. Front-truncated Dome Mode 1487
3. Rear-truncated Dome Mode 1488
4.Environment Map Mode 1488
5. 360Âº Spherical Panoramic 1489

Warp Data Mesh **1490**

Physics Engine **1492**
Description 1492
Options 1492

Bullet physics Python API **1493**

The VideoTexture module **1494**
How it works 1494
Game preparation 1494
First example 1495

Table of Contents

The VideoTexture module
Find material....1495
Create texture....1495
Make texture persistent....1496
Create a source....1496
Setup the source....1497
Play the video....1498
Checking video status....1498
Advanced work flow....1498
Download the demo....1499
Advanced demos....1499

Features....1500

External Resources....1501
External documentation....1501
Community....1501
Development....1501
Game Development tools and utilities....1501
Python....1501

About Blender Wiki PDF Manual

This is an unofficial PDF version of the Blender Wiki Manual pages.

I produced this PDF copy for my personal use since I needed it to learn using Blender, and I could not find an up to date alternative to my knowledge. I've read that Blender documentation is released under the Open Content License (http://opencontent.org/opl.shtml). This license is reported at the end of this pages, below. I am trying to stick to this licence, but i'm no license expert: if you feel something is wrong, just let me know why and possibly how to fix.

Modifications to the wiki content

I did not modify the content of the wiki pages except for removing the parts typical of a web page, which are not needed in a "book", since you don't have to browse through web links, but scroll pages.

Previous versions

In the past i built the wiki manual pdf using different tools to:

- collect links (http://gemal.dk/mozilla/linky.html)
- download pages and editing them (http://www.php.net)
- build the pdf with images (http://www.htmldoc.org)
- build the pdf bookmarks (http://jpdfbookmarks.altervista.org)

trying to automate everything as possibile... I never liked at all that workflow, but it worked in some way.

For this release i managed to use only two tools:

- http://www.php.net

 Used it to scan the base page, find links, download locally pages and images, and build the .book for the next tool
- http://www.htmldoc.org

 Open the .book file and click "generate". After a minute or so, the PDF is ready.

Still, there are some (minor) problems:

- many internal links are broken: some of them work and some do not. I don't know why. When they work, often they jump to the "previous" page they refer to. Try to advance of one page, in case;
- there are strange symbols here and there: the HTMLdoc version I have (1.8.27) is not UTF-8 compatible, it renders web pages using ISO-8859-x encodings. I use ISO-8859-1 since i think it's the best for english Language. New versions or other tools may improve this aspect, in the future. Is not something that makes text difficult to understand, though.
- Size is huge: there are many pages, and many big images and they're useful. I could reduce quality setting for JPGs below 85%, if needed but quality will suffer. Compressing the PDF does not change much the size (tried 7z ultra)

You should be able to know what's changed since the last release browsing here:
http://wiki.blender.org/index.php/Special:Recentchangeslinked/Manual

Where to find this PDF

Up to now, you should be able to find it here: http://www.letworyinteractive.com/b/category/blender/. If any change should happen, i'll post on major forums/blogs for everyone to know.

Since in the past months Nathan had some troubles with his site, check also here
http://www.archive.org/details/BlenderPdfWikiManualUnofficial20090129_630

Alternatives

To my knowledge. See:

1. http://wiki.blender.org/index.php/Meta/Wiki_to_PDF
2. http://www.blender.it/manuali.html (bottom of the page)
3. http://www.blender3d.org/e-shop/product_info.php?products_id=79
4. http://www.blendernation.com/2006/04/20/download-the-blender-manual/
5. Download the entire html structure on your PC with tools like wget or winhttrack

Obviously, there are other books or tutorials or videos/dvd about Blender you can download or buy. They're all useful stuff. Here I am referring only to offline versions of the official Blender manual.

How to contact me

I do this in the spare time so i can't spend too much time on it but if you want to suggest me improvements or other ways to do this, or alternatives, please feel free to do so. Here's how: "m.ardito" is the username and the domain is "libero.it". you know how to use them ;). Please don't post the "reconstructed" address, in no web page, blog, mailing list or newsgroup, anywhere. I already have enough spam to deal with! Thanks.

Have fun! Marco Ardito

Open Content License

OpenContent License (OPL)

Version 1.0, July 14, 1998.

This document outlines the principles underlying the OpenContent (OC) movement and may be redistributed provided it remains unaltered. For legal purposes, this document is the license under which OpenContent is made available for use.

The original version of this document may be found at http://opencontent.org/opl.shtml

LICENSE

Terms and Conditions for Copying, Distributing, and Modifying

Items other than copying, distributing, and modifying the Content with which this license was distributed (such as using, etc.) are outside the scope of this license.

1. You may copy and distribute exact replicas of the OpenContent (OC) as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the OC a copy of this License along with the OC. You may at your option charge a fee for the media and/or handling involved in creating a unique copy of the OC for use offline, you may at your option offer instructional support for the OC in exchange for a fee, or you may at your option offer warranty in exchange for a fee. You may not charge a fee for the OC itself. You may not charge a fee for the sole service of providing access to and/or use of the OC via a network (e.g. the Internet), whether it be via the world wide web, FTP, or any other method.

2. You may modify your copy or copies of the OpenContent or any portion of it, thus forming works based on the Content, and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified content to carry prominent notices stating that you changed it, the exact nature and content of the changes, and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the OC or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License, unless otherwise permitted under applicable Fair Use law.

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the OC, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the OC, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it. Exceptions are made to this requirement to release modified works free of charge under this license only in compliance with Fair Use law where applicable.

3. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to copy, distribute or modify the OC. These actions are prohibited by law if you do not accept this License. Therefore, by distributing or translating the OC, or by deriving works herefrom, you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or translating the OC.

NO WARRANTY

4. BECAUSE THE OPENCONTENT (OC) IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE OC, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE OC "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK OF USE OF THE OC IS WITH YOU. SHOULD THE OC PROVE FAULTY, INACCURATE, OR OTHERWISE UNACCEPTABLE YOU ASSUME THE COST OF ALL NECESSARY REPAIR OR CORRECTION.

5. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MIRROR AND/OR REDISTRIBUTE THE OC AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE OC, EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Theory|Tutorials

Introduction

What is Blender?
- Introduction
- Blender's History
- License
- Blender's Community
- About this Manual

Installing Blender
- Introduction
- Python
- Installing on Windows
- Installing on Linux
- Installing on Mac
- Installing on other Operating Systems
- Configuring Blender

Theory|Tutorials

Starting

The Interface
- Introduction
- Keyboard and Mouse

Your First Animation
- 1/2: A static Gingerbread Man

- Window System
 - Arranging frames
 - Headers
 - Console window
- Window Types
- Screens (Workspace Layouts)
- Scenes
- Configuration
- Contexts
- Menus
- Panels
- Buttons and Controls
- 2/2: Animating the Gingerbread Man

The Vital Functions

- File operations
- Quick render
- Screenshots
- Default scene
- User Preferences
- Undo and Redo
- Help!

Theory|Tutorials

Interaction in 3D

- Introduction

Navigation in 3D Space

- Introduction
- 3D View
- 3D View Options
- 3D View Usage
- Camera View
- Layers
- Local or Global View

Sketch in 3D Space

- Introduction to Grease Pencil
- Drawing sketches
- Layers and Animation
- Converting sketches to geometry

Manipulation in 3D Space

- Introduction
- Hotkeys
- Manipulators
- Manual/Gestures
- Axis Locking
- Pivot Points
- Proportional Edit
- Transform Orientations
- Transform Properties

Theory|Tutorials

Data System

- Blenderâ ← s Library and Data System
- The Datablocks
- Scenes
- Working with Scenes
- The Outliner Window
- Appending and Linking

New Modeling chapter
There is a new version of this chapter pending for revision here, so if you want more up-to-date info, or to edit these pages, please check it!
--Mont29 15:37, 1 April 2010 (UTC)

Theory|Tutorials

Modeling

- Introduction

Objects

- Objects
- Selecting in Object Mode
- Editing Objects
- Groups and Parenting
- Tracking
- Duplication
 - DupliVerts
 - DupliFaces
 - DupliGroup
 - DupliFrames

Meshes

- Meshes
- Mesh Structures
- Mesh Primitives
- Mesh Smoothing
- Selecting Meshes
- Basic Mesh Tools
- Advanced Mesh Tools
- Booleans
- Edge Tools
- Face Tools
- Advanced Edge and Face Tools
- Snap to Mesh
- Vertex Groups
- Weight Paint
- Subdivision Surfaces

Curves

- Curves
- Editing Curves
- Curves
- Deform
- Curves Taper

Surfaces

- Surfaces
- Editing Surfaces
- Surfaces
- Skinning

Text

- Editing Text

Meta Objects

- Meta Objects
- Editing Meta Objects

Scripts

- Modeling Scripts

- Multi Resolution Mesh
- SculptMode
- Retopo

Theory|Tutorials

Modifiers and Deformation

- Introduction
- Modifiers Stack

Deforming (meshes, curves, surfaces, texts, lattices)

- Armature
- Cast
- Curve
- Hooks
- Lattice
- MeshDeform
- Shrinkwrap
- SimpleDeform
- Wave

Constructive (meshes only)

Objects

- Array
- ParticleInstance

Animation

- Build

Mesh

- Bevel
- Booleans
- Decimate
- Displace
- EdgeSplit
- Explode
- Mask
- Mirror
- Smooth
- Subsurf

Textures

- UVProject

Theory|Tutorials

Lighting

- Introduction

Lights

- Introduction
- Light Properties
- Light Attenuation
- Light Textures
- What Light Affects

Radiosity

- Introduction
- Rendering
- Baking

Scene Light

- Ambient
- Light

- Lights In Other Contexts
- Ambient Occlusion
- Exposure

Shadows
- Introduction
- Shadow Properties
- Raytraced Shadow Properties

Volumetric Lights
- Introduction

Lamps
- Introduction
- Lamp Light
 - Raytraced Shadows
- Spot Light
 - Raytraced Shadows
 - Buffered Shadows
 - Halos
- Area Light
 - Raytraced Shadows
- Hemi Light
- Sun Light
 - Raytraced Shadows
 - Sky & Atmosphere
- Lighting Rigs

Introduction to Shading

Theory|Tutorials

Materials

- Introduction

Usage
- Assigning a material
- Material Preview
- Material Options
- Multiple Materials

Properties
- Diffuse Shaders

Node Materials
- Node Materials
- Material Nodes Editor
- Nodes usage
- Material Node Types

Vertex Paint

Specular Shaders
Ambient Light
Effect
Color Ramps
Raytraced Reflections
Raytraced Transparency
Subsurface Scattering (SSS)
Strands

Using Vertex Paint

Halos

Halos

Theory|Tutorials

Textures

Introduction

Options

Texture Options
Texture Channels
Map Input
Map To

Texture maps

Bump and Normal Maps
Displacement Maps

Texture Types

Procedural Textures
Image Textures
Video Textures
Environment Maps
Texture nodes
Texture Plugins

UV Textures

UV Unwrapping Explained
Unwrapping a Mesh
Managing the UV Layout
Editing the UV Layout
Applying an Image
Painting the Texture

Projection Paint

Theory|Tutorials

World and Ambient Effects

Introduction
World
Background
Mist
Stars

Working on Animation chapters
Please avoid modifying the four next chapters about animation, Iâ← m starting an update of the whole section â← see this pageâ←¦
--Mont29 13:14, 27 April 2010 (UTC)

Theory|Tutorials

Animation Basics

Introduction

Tools
Ipo Types
Creating Ipo Keyframes
Editing Ipo Curves and Keyframes
The Timeline

Animating Objects
Moving objects on a Path
Changing Object Layers

Animating Deformation
Methods of deformation
Shape Keys
Absolute Shape Keys
Deforming by a Lattice

Other Resources
Reference: Editing Ipos
Hook Modifier to animate mesh deformation

BSoD Tutorial
Constraints and Axis Locks

Working on Animation chapters
The new rigging chapter is now written, and in proofreading state â← Â click here to access its index. Even though it might still contain things to correct/enhance, it is (I thinkâ←¦) much more precise and up to date than the current one belowâ←¦
--Mont29 12:23, 1 June 2010 (UTC)

Theory|Tutorials

Armatures and Character Animation

IntroductionMesh Skin Weighting

Armatures
- Armature Objects
- Editing Armatures
- Posing Armatures
- Inverse Kinematics
- Mesh Skin Weighting

Animation Editors
- The Action Editor
- Non Linear Animation

Theory|Tutorials

Constraints

Introduction
- Introduction
- Constraint Stack

Objects Constraints
- Child Of
- Copy Location
- Copy Rotation
- Copy Scale
- Limit Location
- Limit Rotation
- Limit Scale
- Floor
- Track To
- Locked Track
- Follow Path
- Clamp To
- Stretch To
- Transformation

Special Constraints
- Rigid Body Joint
- IK Solver
- Action
- Script
- Null

Theory|Tutorials

Advanced Animation

Driven Shape Keys (Ipo Drivers)
Ipo Drivers Example
Python
Ipo
Drivers

Theory|Tutorials

Effects and Physical Simulation

Introduction

Particles
- Particles
- Types
- Physics
- Visualization
- Controlling Emission, Interaction and Time
- Children
- Particle Mode
- Hair

Soft Body
- Introduction
- Exterior forces
- Interior forces
- Collisions
- Simple Examples
- Combination with armatures
- Combination with hair particles
- Reference

Force Fields
Collision
Rigid Bodies
Fluids
Clothes

Theory|Tutorials

Rendering

Introduction
The Render Window

Options
- Render Options

Compositing
- Render Layers
- Render Passes

Camera

Oversampling (Antialiasing)
Render engines: YafRay
Render Bake
Rendering From the Command Line
Perspective (Vanishing points)
Depth Of Field

Output
Output Options
Output Formats
Rendering Animations
Video output

Other
Rendering Performances

Post-production

Theory|Tutorials

Compositing with nodes

Compositing Nodes
Introduction
Nodes Editor
Using Nodes

Node types
Node types
Input Nodes
Output Nodes
Color Nodes
Vector Nodes
Filter Nodes
Convertor Nodes
Matte Nodes
Distortion Nodes

Theory|Tutorials

Editing Sequences

Introduction
The sequencer

Effects
Built-in Effects

- Sequencer window modes
- Sequence Screen Layout
- Usage
- Plugin Effects
- Audio
- Audio Sequences

Theory|Tutorials

Extending Blender

- Introduction
- Python Scripting
 - Python Scripting in Blender
 - Setting up Python
 - The Text Editor
 - A working example
 - References
- Plugins
 - Blender's Plugins System
- Texture
 - Texture plugins specifications
- Sequence
 - Sequence plugins specifications

Python Scripts
- Bundled Scripts
- Catalog of all Scripts

Blender Python API
- Blender2.49 Python API

Theory|Tutorials

Game Engine

- Introduction
- Physics
 - Physics Engine
- Usage
 - Logic Panel
 - Python API
 - Bullet physics
 - VideoTexture

- Object type
 - Occluder
 - Various collision
 - Sensor
 - Static
 - Dynamic
 - Rigid body
 - Soft body
- Properties
- Sensors
 - Sensor types
- Controllers
- Actuators
 - Motion
 - Shape Action
 - Action
 - IPO
 - Property
 - Edit Object
 - Scene
 - 2D Filters
- Camera
 - Dome camera

- Blender2.49 Game Engine Python API
- No
- Game Engine Basics (BSoD Tutorial)
- FAQ and List of Features
- External resources

Introduction

Welcome to Blender! The documentation of Blender consists of many parts: this user manual, a reference guide, tutorials, forums, and many other web resources.

The first part of this manual will guide you through downloading Blender executable and installing it, OR downloading the source files and then building an executable file that runs on your machine.

Blender has a very unusual interface, highly optimized for 3D graphics production. This might be a bit confusing to a new user, but will prove its strength in the long run.

We highly recommend you to read our section on The Interface carefully, both to get familiar with the interface and with the conventions used in the documentation.

What is Blender?

Blender was first conceived in December 1993 and born as a usable product in August 1994 as an integrated application that allows the creation of a broad range of 2D and 3D content. Blender provides a broad spectrum of modeling, texturing, lighting, animation and video post-processing functionality in one package. Through its open architecture, Blender provides cross-platform interoperability, extensibility, an incredibly small footprint, and a tightly integrated workflow. Blender is one of the most popular Open Source 3D graphics applications in the world.

Aimed world-wide at media professionals and artists, Blender can be used to create 3D visualizations, stills as well as broadcast and cinema quality videos, while the incorporation of a real-time 3D engine allows for the creation of 3D interactive content for stand-alone playback.

Originally developed by the company 'Not a Number' (NaN), Blender now is continued as 'Free Software', with the source code available under the GNU GPL license. Development is now conducted under the leadership of the Blender Foundation in the Netherlands.

Key Features:

- Fully integrated creation suite, offering a broad range of essential tools for the creation of 3D content, including modeling, uv-mapping, texturing, rigging, skinning, animation, particle and other simulation, scripting, rendering, compositing, post-production, and game creation;
- Cross platform, with OpenGL uniform GUI on all platforms, ready to use for all versions of Windows (98, NT, 2000, XP, Vista, 7), Linux, OS X, FreeBSD, Irix, Sun and numerous other operating systems;
- High quality 3D architecture enabling fast and efficient creation work-flow;
- More than 200,000 downloads of each release (users) worldwide;
- User community support by forums for questions, answers, and critique at http://BlenderArtists.org and news services at http://BlenderNation.com;
- Small executable size, easy distribution;

You can download the latest version of Blender here.

Blenderâ← s History

In 1988 Ton Roosendaal co-founded the Dutch animation studio *NeoGeo*. NeoGeo quickly became the largest 3D animation studio in the Netherlands and one of the leading animation houses in Europe. NeoGeo created award-winning productions (European Corporate Video Awards 1993 and 1995) for large corporate clients such as multi-national electronics company Philips. Within NeoGeo Ton was responsible for both art direction and internal software development. After careful deliberation Ton decided that the current in-house 3D tool set for NeoGeo was too old and cumbersome to maintain and upgrade and needed to be rewritten from scratch. In 1995 this rewrite began and was destined to become the 3D software creation suite we all now know as *Blender*. As NeoGeo continued to refine and improve Blender it became apparent to Ton that Blender could be used as a tool for other artists outside of NeoGeo.

In 1998, Ton decided to found a new company called Not a Number (NaN) as a spin-off of NeoGeo to further market and develop Blender. At the core of NaN was a desire to create and distribute a compact, cross platform 3D creation suite for free. At the time this was a revolutionary concept as most commercial modelers cost several thousands of (US) dollars. NaN hoped to bring professional level 3D modeling and animation tools within the reach of the general computing public. NaNâ← s business model involved providing commercial products and services around Blender. In 1999 NaN attended its first Siggraph conference in an effort to more widely promote Blender. Blenderâ← s first 1999 Siggraph convention was a huge success and gathered a tremendous amount of interest from both the press and attendees. Blender was a hit and its huge potential confirmed!

On the wings of a successful Siggraph in early 2000, NaN secured financing of â↓¬4.5m from venture capitalists. This large inflow of cash enabled NaN to rapidly expand its operations. Soon NaN boasted as many as fifty employees working around the world trying to improve and promote Blender. In the summer of 2000, Blender v2.0 was released. This version of Blender added the integration of a game engine to the 3D suite. By the end of 2000, the number of users registered on the NaN website surpassed 250,000.

Unfortunately, NaNâ← s ambitions and opportunities didnâ← t match the companyâ← s capabilities and the market realities of the time. This over-extension resulted in restarting NaN with new investor funding and a smaller company in April 2001. Six months later NaNâ← s first commercial software product, *Blender Publisher* was launched. This product was targeted at the emerging market of interactive web-based 3D media. Due to disappointing sales and the ongoing difficult economic climate, the new investors decided to shut down all NaN operations. The shutdown also included discontinuing the development of Blender. Although there were clearly shortcomings in the then current version of Blender, with a complex internal software architecture, unfinished features and a non-standard way of providing the GUI, with the enthusiastic support from the user community and customers who had purchased Blender Publisher in the past, Ton couldnâ← t justify leaving Blender to disappear into oblivion. Since restarting a company with a sufficiently large team of developers wasnâ← t feasible, in March 2002 Ton Roosendaal founded the non-profit organization *Blender Foundation*.

The Blender Foundationâ← s primary goal was to find a way to continue developing and promoting Blender as a community-based Open Source project. In July 2002, Ton managed to get the NaN investors to agree to a unique Blender Foundation plan to attempt to release Blender as open source. The â← Free Blenderâ← campaign sought to raise â↓¬100,000 so that the Foundation could buy the rights to the Blender source code and intellectual property rights from the NaN investors and subsequently release Blender to the open source community. With an enthusiastic group of volunteers, among them several ex-NaN employees, a fund raising campaign was launched to â← Free Blenderâ← . To everyoneâ← s surprise and delight the campaign reached the â↓¬100,000 goal in only seven short weeks. On Sunday October 13, 2002, Blender was released to the world under the terms of the GNU General Public License (GPL). Blender development continues to this day driven by a team of far-flung, dedicated volunteers from around the world led by Blenderâ← s original creator, Ton Roosendaal.

Version/Revision Milestones

Blenderâ← s history and road-map:

- 1.00 â← January 1995: Blender in development at animation studio NeoGeo.
- 1.23 â← January 1998: SGI version published on the web, IrisGL.
- 1.30 â← April 1998: Linux and FreeBSD version, port to OpenGL and X.
- 1.3x â← June 1998: NaN founded.
- 1.4x â← September 1998: Sun and Linux Alpha version released.
- 1.50 â← November 1998: First Manual published.
- 1.60 â← April 1999: C-key (new features behind a lock, $95), Windows version released.
- 1.6x â← June 1999: BeOS and PPC version released.
- 1.80 â← June 2000: End of C-key, Blender full freeware again.
- 2.00 â← August 2000: Interactive 3D and real-time engine.
- 2.10 â← December 2000: New engine, physics, and Python.
- 2.20 â← August 2001: Character animation system.
- 2.21 â← October 2001: Blender Publisher launch.
- 2.2x â← December 2001: Mac OSX version.
- **13 October 2002: Blender goes Open Source, 1st Blender Conference**.
- 2.25 â← October 2002: Blender Publisher becomes freely available.
- Tuhopuu1 â← Oct 2002: The experimental tree of Blender is created, a coderâ← s playground.
- 2.26 â← February 2003: The first true Open Source Blender.
- 2.27 â← May 2003: The second Open Source Blender.
- 2.28x â← July 2003: First of the 2.28x series.
- 2.30 â← October 2003: Preview release of the 2.3x UI makeover presented at the 2nd Blender Conference.
- 2.31 â← December 2003: Upgrade to stable 2.3x UI project.
- 2.32 â← January 2004: Major overhaul of internal rendering capabilities.
- 2.33 â← April 2004: Game Engine returns, ambient occlusion, new procedural textures.
- 2.34 â← August 2004: Big improvements: particle interactions, LSCM UV mapping, functional YafRay integration, weighted creases in subdivision surfaces, ramp shaders, full OSA, and many many more.
- 2.35 â← November 2004: Another version full of improvements: object hooks, curve deforms and curve tapers, particle duplicators and much more.
- 2.36 â← December 2004: A stabilization version, much work behind the scene, normal and displacement mapping improvements.
- 2.37 â← June 2005: A big leap: transformation tools and widgets, softbodies, force fields, deflections, incremental subdivision surfaces, transparent shadows, and multithreaded rendering.
- 2.40 â← December 2005: An even bigger leap: full rework of armature system, shape keys, fur with particles, fluids and rigid bodies.
- 2.41 â← January 2006: Lots of fixes, and some game engine features.
- 2.42 â← July 2006: The **Node** release. Over 50 developers contributed nodes, array modifier, vector blur, new physics engine, rendering, lipsync and, many other features. This was the release following Project Orange.
- 2.43 â← February 2007: The **Multi** release: multi-resolution meshes, multi-layer UV textures, multi-layer images and multi-pass rendering and baking, sculpting, retopology, multiple additional matte, distort and filter nodes, modeling and animation improvements, better painting with multiple brushes, fluid particles, proxy objects, sequencer rewrite, and post-production UV texturing. whew! Oh, and a website rewrite. And yes, it still has multi-threaded rendering for multi-core CPUs. With Verse it is multi-user, allowing multiple artists to work on the same scene collaboratively. Lastly, render farms still provide multi-workstation distributed rendering.
- 2.44 â← May 2007: The **SSS** release: the big news, in addition to two new modifiers and re-awakening the 64-bit OS support, was the addition of subsurface scattering, which simulates light scattering beneath the surface of organic and soft objects.

- 2.45 â← September 2007: Another bugfix release: serious bugfixes, with some performance issues addressed.
- 2.46 â← May 2008: The **Peach** release was the result of a huge effort of over 70 developers providing enhancements to the core and patches to provide hair and fur, a new particle system, enhanced image browsing, cloth, a seamless and non-intrusive physics cache, rendering improvements in reflections, AO, and render baking; a mesh deform modifier for muscles and such, better animation support via armature tools and drawing, skinning, constraints and a colorful Action Editor, and much more. It was the release following Project Peach.
- 2.47 â← August 2008: Bugfix release.
- 2.48 â← October 2008: The **Apricot** release: cool GLSL shaders, lights and GE improvements, snap, sky simulator, shrinkwrap modifier, python editing improvements.
- 2.49 â← June 2009: The **Pre-Re-Factor** release added significant enhancements to the core and GE. Core enhancements include node-based textures, armature sketching (called Etch-a-Ton), boolean mesh operation improvements, JPEG2000 support, projection painting for direct transfer of images to models, and a significant Python script catalog. GE enhancements included video textures, where you can play movies in-game (!), upgrades to the Bullet physics engine, dome (fish-eye) rendering, and more API GE calls made available.

About Free Software and the GPL

When one hears about "free software", the first thing that comes to mind might be "no cost". While this is true in most cases, the term "free software" as used by the Free Software Foundation (originators of the GNU Project and creators of the GNU General Public License) is intended to mean "free as in freedom" rather than the "no cost" sense (which is usually referred to as "free as in free beer"). Free software in this sense is software which you are free to use, copy, modify, redistribute, with no limit. Contrast this with the licensing of most commercial software packages, where you are allowed to load the software on a single computer, are allowed to make no copies, and never see the source code. Free software allows incredible freedom to the end user; in addition, since the source code is available universally, there are many more chances for bugs to be caught and fixed.

When a program is licensed under the GNU General Public License (the GPL):

- you have the right to use the program for any purpose;
- you have the right to modify the program, and have access to the source codes;
- you have the right to copy and distribute the program;
- you have the right to improve the program, and release your own versions.

In return for these rights, you have some responsibilities if you distribute a GPL'd program, responsibilities that are designed to protect your freedoms and the freedoms of others:

- You must provide a copy of the GPL with the program, so that the recipient is aware of his rights under the license.
- You must include the source code or make the source code freely available.
- If you modify the code and distribute the modified version, you must license your modifications under the GPL and make the source code of your changes available. (You may not use GPL'd code as part of a proprietary program.)
- You may not restrict the licensing of the program beyond the terms of the GPL. (You may not turn a GPL'd program into a proprietary product.)

For more on the GPL, check the GNU Project Web site. For reference, a copy of the GNU General Public License is included in Volume II.

Getting support: the Blender community

Being freely available from the start, even while closed source, helped a lot in Blender's adoption. A large, stable and active community of users has gathered around Blender since 1998.

The community showed its best in the crucial moment of freeing Blender itself, going Open Source under GNU GPL in late summer 2002.

The community itself is now subdivided into two, widely overlapping sites:

1. The Development Community, centered around the Blender Foundation site. Here you will find the home of the development projects, the Functionality and Documentation Boards, the CVS repository with Blender sources, all documentation sources, and related public discussion forums. Developers coding on Blender itself, Python scripters, documentation writers, and anyone working for Blender development in general can be found here.
2. The User Community, centered around the independent site BlenderArtists. Here Blender artists, Blender gamemakers and Blender fans gather to show their creations, get feedback on them, and ask for help to get a better insight into Blender's functionality. Blender Tutorials and the Knowledge Base can be found here as well.

These two websites are not the only Blender resources. The Worldwide community has created a lot of independent sites, in local languages or devoted to specialized topics. A constantly updated listing of Blender resources can be found at the above mentioned sites.

For immediate online feedback there are three IRC chat channels permanently open on irc.freenode.net. You can join these with your favorite IRC client.

The IRC channels are #blenderchat for general discussion of blender; #blenderqa for asking questions on Blender usage; and #gameblender for discussion on issues related to game creation with Blenders included game engine. For developers there is also #blendercoders for developers to ask questions and discuss development issues, as well as a meeting each Sunday at 4 pm Netherlands time; #blenderpython for discussion of the python API and script development; #blenderwiki for questions related to editing the wiki

Who uses Blender?

New releases of Blender are downloaded by more than a million people around the world just in the first 10 days of release. This figure spans all platforms (Windows, Linux, and MacOS) and does not include redistribution, which is fully authorized and unrestricted. We estimate there are in excess of two million users. This manual is written to serve the wide array of talented people that use Blender:

- Hobbyist/Student that just wants to explore the world of computer graphics (CG) and 3D animation
- 2-D artist that produces single image art/posters or enhances single images as an image post-processing lab
- 2-D artist or team that produces cartoon/caricature animations for television commercials or shorts (such as â← The Magic of Ameliaâ←)
- 3-D artist that works alone or with another person to produce short CG animations, possibly featuring some live action (such as "Suburban Plight").
- 3-D team that produces an animated (100% CG) movie (such as "Elephant's Dream", "Plumiferos").
- 3-D team that works together to produce live action movies that include some CG.

A wide range of age groups, from teenagers to oldsters use Blender, and the user community is fairly evenly divided between novice and professional graphic artists; those occasional users as well as commercial houses. We can divide the 2-D and 3-D teams that produce movies and animations further into individual job categories. Those that use Blender include:

- Director - Defines what each Scene should contain, and the action (animation) that needs to occur within that scene. Defines shots (camera takes) within that scene.
- Modeler - Makes a virtual reality. Specialties include Character, Prop and Landscapes/Stage modelers
- Cameraman, Director of Photography (DP): sets up the camera and its motion, shoots the live action, renders the output frames.
- Material Painter - paints the set, the actors, and anything that moves. If it doesn't move, they paint it anyway.
- Animation and Rigging - makes things hop about using armatures
- Lighting and Color Specialist - Lights the stage and sets, adjusts colors to look good in the light, adds dust and dirt to materials, scenes, and textures.
- Special Purpose talent - Fluids, Motion Capture, Cloth, dust, dirt, fire, explosions, you know, the fun stuff
- Editor - takes all the raw footage from the DP and sequences it into an enjoyable movie. Cuts out unnecessary stuff.

About this Manual

This manual is a mediawiki implementation that is written by a world-wide collaboration of volunteer authors. It is updated daily, and this is the English version. Other language versions are translated, generally, from this English source for the convenience of our world-wide audience. It is constantly out of date, thanks to the tireless work of some 50 or more volunteer developers, working from around the world on this code base. However, it is the constructive goal to provide you with the best possible professional documentation on this incredible package.

To assist you in the best and most efficient way possible, this manual is organized according to the creative process generally followed by 3D artists, with appropriate stops along the way to let you know how to navigate your way in this strange territory with a new and deceptively complex software package. If you read the manual linearly, you will follow the path most artists use in both learning Blender *and* developing fully animated productions:

1. Getting to know Blender = Intro, Navigating in 3D, scene management
2. Models = Modeling, Modifiers
3. Lighting
4. Shading = Materials, Textures, Painting, Worlds & Backgrounds
5. Animation = Basics, Characters, Advanced, Effects & Physical Simulation
6. Rendering = Rendering, Compositing, Video Sequence Editing
7. Beyond Blender = Extending Blender
8. Game Engine = Interactive 3D and Gaming



Expandable Images in Manual:

Usually throughout this manual, if there is a little expanding icon in the bottom right-hand corner of the image, you may click it to see the image in a larger format. Also if when moving your mouse pointer over the image it changes shape to show the image can be clicked, this usually also means the image can be expanded/enlarged. Be aware that different wiki theme layouts can alter the location and appearance of the expand image indicator.

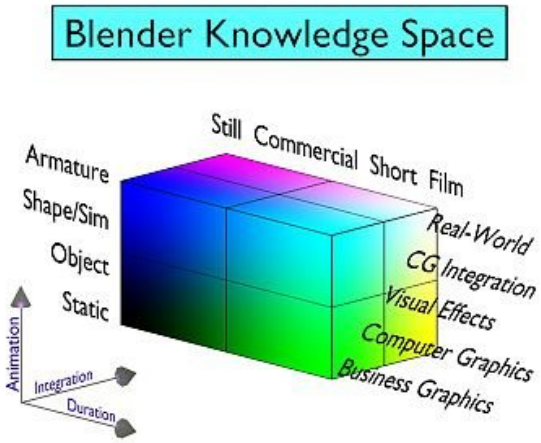
Audience

This manual is written for a very broad audience, to answer the question "I want to *do something*; how do I do it using Blender?" all the way to "what is the latest change in the way to sculpt a mesh?"

This manual is a worldwide collaborative effort using time donated to the cause. While there may be some lag between key features being implemented and their documentation, we do strive to keep it as up-to-date as possible. We try to keep it narrowly focused on what you, the end user, need to know, and not digress too far off topic, as in discussing the meaning of life.

There are other Blender wiki books that delve deeper into other topics and present Blender from different viewpoints, such as the Tutorials, the Reference Manual, the software itself, and its scripting language. So, if a question is not answered for you in this User Manual, please search the other Blender wiki books.

Learning CG and Blender



Getting to know Blender and learning Computer Graphics (CG) are two different topics. On the one hand, learning what a computer model is, and then learning how to develop one in Blender are two different things to learn. Learning good lighting techniques, and then learning about the different kinds of lamps in Blender are two different topics. The first, or conceptual understanding, is learned by taking secondary and college courses in art and media, by reading books available from the library or bookstore on art and computer graphics, and by trial and error. Even though a book or article may use a different package (like Max or Maya) as its tool, it may still be valuable because it conveys the concept.

Once you have the conceptual knowledge, you can easily learn Blender (or any other CG package). Learning both at the same time is difficult, since you are dealing with two issues. The reason for mentioning this is to make you aware of this dilemma, and how this manual attempts to address both topics in one wiki book. The conceptual knowledge is usually addressed in a short paragraph or two at the beginning of a topic or chapter, that explains the topic and provides a workflow, or process, for accomplishing the task. The rest of the manual section addresses the specific capabilities and features of Blender. The user manual cannot give you the full conceptual knowledge - that comes from reading books, magazines, tutorials and sometimes a life-time of effort. You can use Blender to produce a full-length feature film, but reading this manual and using Blender won't, in itself, make you another Steven Spielberg!

At a very high level, using Blender can be thought of as knowing how to accomplish imagery within three dimensions or phases of activity:

1. Integration - rendering computer graphics, working with real-world video, or mixing the two (CGI and VFX)
2. Animation - posing and making things change shape, either manually or using simulation
3. Duration - producing a still image, a short video, a minute-long commercial, a ten minute indie short, or a full-length feature film.

Skills, like navigating in 3D space, modeling, lighting, shading, compositing, and so forth are needed to be productive in any given area within the space. Proficiency in a skill makes you productive. Tools within Blender have applicability within the space as well. For example, the video sequence editor (VSE) has very little to do with the skill of animation, but is deeply applicable along the Duration and Integration scales. From a skills-learning integration perspective, it is interesting to note that the animation curve, called an Ipo curve, is used in the VSE to animate effects strips.

Most people's interest's lie at any given time at the corners/intersections depicted in the Blender Knowledge Space, which are destinations, if you will. For example, there are many talented artists who produce Static-Still-CG images. Tony Mullen's book *Introducing Character Animation With Blender* pertains to using CG models deformed by Armatures and shapes to produce short animations. Using Blender fluids in a TV

production or commercial is at the Shape/Sim-Integrated-Minute intersection. Elephants Dream and Big Buck Bunny are found at the Armature-CG-Indie space. Therefore, depending on what you want to do, various tools and topics within Blender will be of more or less interest to you.

A fourth dimension is Game Design, which incorporates all of this knowledge with Gaming wrapped around it as well. A game not only has a one-minute cinematic in it, but it also has actual game play, story line programming, etc. -- which may explain why it is so hard to make a game; you have to understand all this stuff before you can actually construct a game. Therefore, this Manual does not address using the Game Engine; that is a whole 'nother wiki book.

Installing the Binaries

Blender is available both as a binary executable and as source code on the Foundation site (http://www.blender.org/). At the main page click on the 'Downloads' section.

For the online manual hosted at the wiki, you can generally use the most recent version of Blender located at the Blender Foundation website (although all of the features from the newest release version may not be fully updated). If you are using a published version of this manual it is recommended that you use the Blender version included on the Guide CD-ROM. In the following text, whenever "download" is mentioned, those using the book should instead retrieve Blender from the CD-ROM.

Downloading and installing the binary distribution

The binary distributions are provided for primary operating system families. Some unofficial distributions may exist for other operating systems, but they're not supported by the Blender Foundation and you should report directly to their maintainers:

- Windows
- Linux
- MacOSX
- FreeBSD, Irix, Solaris
- MorphOS (not supported by Blender Foundation)

OpenGL is used in Blender for all drawing, including menus and buttons. For this, you will need OpenGL installed on your system. This dependency makes a proper and compliant OpenGL installation on your system a requirement. Generally speaking integrated graphics chips and older low end graphics cards will perform poorly or not at all with Blender due to their often poor support for OpenGL. (It is often possible to work around the poor OpenGL support of such cards by using software-based OpenGL solutions such as by turning down or off hardware acceleration on Windows, or using software MESA 3D on Linux). Running Blender using the blender-softwaregl script in the distribution may help in this case.

Rendering is done by the Blender rendering engine in core memory and by the main CPU of your machine, so an unsupported graphics card will not have an impact if you use the machine only for rendering (as would be the case for a rendering farm).

The installer will create files and several folders in two locations on your computer: one set of folders is for Blender itself, and the other is a user set of folders for your user data. You must have administrator authorization to create these. The folders are:

- .blender - configuration information (mostly prompts in your native language)
- blendcache_.B - temporary space for physics simulation information (softbodies, cloth, fluids)
- plugins - added functionality for textures and sequencing
- scripts - python scripts that extend Blender functionality
- tmp - temporary output, intermediate renders

Hardware Support

Blender supports 64-bit hardware platforms running a 64-bit unix operating system, removing the 2Gig addressable memory limit. Work is underway to support a Windows 64-bit OS (call for developer help!)

Blender supports multi-CPU chips, like the Intel Core-Duo and AMD X2 chips by providing a *Threads:* setting when rendering to work both cores in parallel when rendering an image.

Blender supports a wide variety of pen-based tablets on all major operating systems, in particular OS X, Windows XP, and Linux OSes.

Tips on making Blender run faster and render swifter can be found here.

Developers platforms

This is the list of systems in use and supported by active Blender developers:

Name	**OS**	**CPU**	**Graphics card**
Andrea Weikert	Windows XP 32	AMD Athlon 64 X2	Nvidia Quadro FX1500
Andrea Weikert	Windows XP 32	Intel P4	ATI Radeon 9000
Benoit Bolsee	Windows XP 32	AMD Athlon XP	ATI Radeon 9200
Brecht van Lommel	Windows Vista 32	Intel Core 2 Duo	NVidia GeForce 9600 GT
Daniel Genrich	Windows Vista 64	Intel Core 2 Duo	NVidia GeForce 8500 GT
Joshua Leung	Windows Vista 32	Intel Core2 Duo	Nvidia GeForce Go 7600
Nathan Letwory	Windows Vista Home Premium 32	AMD Athlon 64 Dual Core	ATI HD 3400
Nathan Letwory	Windows Vista Home Premium 32	AMD Phenom Quad Core	ATI HD 2600 Pro
Nathan Letwory	Windows Vista Home Premium 32	Intel Celeron M	ATI Radeon XPRESS 200M
Robin Allen	Windows XP 32	Intel Centrino duo	NVidia GeForce go 7600
Thomas Dinges	Windows Vista 32	Intel Core 2 Quad	NVidia GeForce 8600 GT
Andrea Weikert	Linux 32	AMD Athlon 64 X2	Nvidia Quadro FX1500
Brecht van Lommel	Linux 32-64	Intel Core 2 Duo	Nvidia GeForce 8800 GTS

Campbell Barton	Linux 32	AMD Athlon 64 X2	Nvidia GeForce 8800 GT
Diego Borghetti	Linux 32	Intel Core 2 Duo	Nvidia 8400 GT
Diego Borghetti	Linux 64	AMD Athlon 64 X2	Nvidia 8600 GT
Ken Hughes	Linux 32	Intel Core Duo	Nvidia GeForce GO 7500
Ken Hughes	Linux 64	AMD Athlon 64 X2	Nvidia GeForce 6600
Kent Mein	Linux 64	Intel Core Duo	Nvidia Quadro FX 1400
Michael Fox	Linux 32	Celeron	Nividia GeForce 6200
Raul Fernandez Hernandez	Linux 32	Pentium D 945	ATI X1550
Robin Allen	Linux 32	Intel Centrino duo	NVidia GeForce go 7600
Brecht van Lommel	OSX 10.4	PPC G4	ATI radeon 9600 (Mobility 9700)
Dustin Martin	OSX 10.5	Dual Quad Intel	Nvidia Geforce 8800 GT
Jean-Luc PeuriÃ¨re	OSX 10.4	PPC dual G5	ATI 9650
Ton Roosendaal	OSX 10.3	PPC dual G5	ATI 9600
Ton Roosendaal	OSX 10.4	Dual Core Intel	ATI x1600
Matt Ebb	OSX 10.5	Dual Core Intel MBP	nVidia 8600M
Kent Mein	SunOS 5.8	Sun Blade 150	ATI PGX
Stefan Gartner	SGI Irix 6.5 (gcc)	R12000	V8
Timothy Baldridge	SGI Irix 6.5 (mipspro)	8 x R16000	(headless)
Timothy Baldridge	SGI Irix 6.5 (mipspro)	2 x R10000	
Marcel	Windows 7 64bit	AMD Athlon 4850e	ATI Radeon HD3200 IGP

Compiling the Source

There are presently four build systems for making a binary for the different operating systems supported. For more information about compiling a custom installation binary for your machine, please see this blender.org page or check this wiki page.

Compiling the Plugins

Plugins are dynamically loaded routines that augment functionality in either texture generation or sequencing (image manipution). See this thread for more information.

Python, the Scripting Language

Python is a general purpose scripting language and there is a special interface to access all of Blender's internal functions from that language. Scripts are written in this language that extend the functionality of Blender, without having to re-compile and link the binary distribution. These scripts are written by user-programmers. The recommended version of Python is normally included and installed with the distribution, however you may also download and install it directly from the official Python website, and install it separately. (When downloading Python separately, do pay attention to the version number, as the greater version number will need to match the version number that Blender was compiled on. The python version will be displayed in the console window when Blender starts. The supported python version for Blender as of the date for this edit will be Python 2.6) Most functions do not rely on Python; a notable exception is the Help menu which opens a web browser pointed to a specific location. Help text is not bundled into Blender; you must download the latest wiki or pdf user manuals, found here or at www.blender.org.

In general, wherever you install Python, you need to establish an operating system environment variable PYTHONPATH and point it to the Blender Scripts directory where python modules are installed, e.g. "C:\Program Files\Blender Foundation\Blender\scripts\bpymodules" for Windows machines. Environment Variables on Windows machines are set in the advanced Systems settings in the Control Panel.

When Blender starts on a machine with Python properly installed, you will see this message in the console window:

```
 Compiled with Python version 2.5.
 Checking for installed Python... got it!
```

The above messages means that you have installed Python and have the full development and execution environment, and will be able to access, execute and run all Python scripts that are bundled or available for Blender. If you see a different message, such as:

```
Could not find platform independent libraries <prefix>
Could not find platform dependent libraries <exec_prefix>
Consider setting $PYTHONHOME to <prefix>[:<exec_prefix>]
'import site' failed; use -v for traceback
Checking for installed Python... No installed Python found.
Only built-in modules are available.  Some scripts may not run.
Continuing happily.
```

it just means that the full Python is not available. If you want full Python functionality, refer to the Python site for installation instructions.

When you install Blender, you must tell the Python module where you put the scripts. if you choose to put user data in a different location for each user, then the install will put your scripts in the

```
'C:\Users\<Current User>\AppData\Roaming\Blender Foundation\Blender\.blender\scripts'
```

folder. If you are upgrading, you probably want to overwrite all your old scripts with the new versions, and not have several versions of the same script hanging around on your PC. The best place, if you will not be editing them, is to put them in your Program Files folder with Blender:

1. Do a search on your machine for a file name with the words 'Scripts'.
2. you will see the scripts folder appear after the initial search....C:\Program Files\Blender-2.46/.blender/scripts or something similar....
3. open the script folder from the search window. you will see all the scripts. You can leave em there or put them on your desktop temporarily....
4. Then go to program files, then to blender foundation, then blender folder, then make a new folder

called scripts in the blender folder....

5. Drag and drop or copy all the scripts from where ever you put them into this folder.
6. Make sure to include the 2 module folders in the script file.
7. Then, if you don't know this already, Open Blender
8. In Blender, the top menu bar hides all the preferences. Drag it down and then you will see a button marked **file paths**.
9. Once you click that File Paths button a set of path fields will be revealed.
10. Go to the script one and drill down to the script folder you just created in blender where you put all the scripts.
11. Then hit the button that says 'Select Script Paths'.
12. Then go to the file menu and save as default setting so Blender will remember that the script folder is where you told it to look CtrlU
13. Be careful though if you have already done stuff in blender at this point every time you start it it will be the default start up.

Windows

Quick Install

Download the file `blender-2.##-windows.exe`, (`2.##` being the version number) from the downloads section of the Blender Website. Start the installation by double-clicking the file. This presents you with some questions, for which the defaults should be OK. After setup is complete, you can start Blender right away, or use the entry in the Start menu.

In-depth Instructions

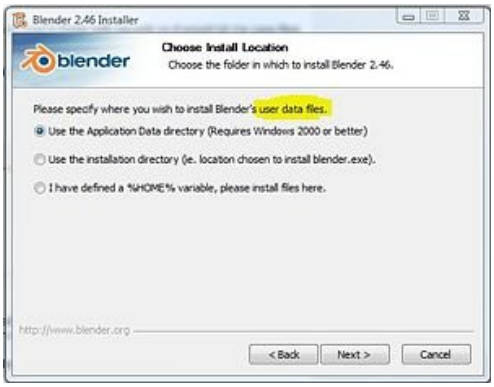


Download the file `blender-2.##-windows.exe` from the 'Downloads' section of the Blender Website. Choose to download it (if prompted), select a location and click "Save". Then navigate to the location you saved the file in and double-click it to start the installation. The first dialog presents you with a helpful reminder: You must be logged on to your PC as ADMINISTRATOR! You will be denied access if you do not have the rights (Vista user especially) to the C:/Program Files folder and/or are authorized to install executable software on the PC. I have Vista install as Administrator (in the file explorer, right-click on the downloaded 'blender-2.46-windows.exe' and select 'Run as administrator' from the right-click menu).

The second dialog presents you with the license. You are expected to accept it if you want the installation to go any further. After accepting the license, select the components you wish to install (there is just one, Blender) and the additional actions:

- Add a shortcut to the Start menu,
- Add Blender's icon to desktop,
- associate .blend files with Blender.

By default they are all checked. If you don't want some action to be taken, simply uncheck it. When done, click on Next. The next dialog is where to put the executable files, usually in the C:\Program Files folder.



The next dialog is tricky, and it is where to put user files. These folders save user data, namely temp data like test renders and physics data. Each user of that PC can have their own, or they can all share one.

Select a place to install the files to (the default should be OK), and click Next to install Blender. Press Close when installation is over.

Afterwards you will be asked whether you want to start Blender immediately. Blender is now installed and can be started by means of the Start menu (an entry named "Blender Foundation" will have been created by the setup routine) or by double-clicking a Blender file (`*.blend`).

After the files are unpacked, Blender will check for required system components, like DLLs, which you must get from Microsoft or your hardware vendor. Most common is a VCRT dll that is/was not bundled with old versions of Windows. After confirmation, you will be able to run Blender!

Portable Install

If, like many people, you are a) obsessed with Blender, and b) have a USB drive, you'll be glad to know that Blender runs beautifully off a USB key. Just download the .zip version and extract it. You may want to avoid having it store the animation output or other temporary files on the drive, as it may shorten the life, but otherwise, Blender runs fine.

Linux

Binaries for the Linux operating systems are provided for two different hardware architectures x86 (Intel and AMD processors) and PowerPC, and the choice between statically linked or dynamically loaded libraries.

The difference between the dynamic and the static binaries is important. The static binaries have the OpenGL libraries compiled in. This makes Blender run on your system without using hardware accelerated graphics. Use the static version if the dynamic version fails!

Quick Install

Download the file `blender-2.##-linux-glibc#.#.#-ARCH.tar.gz` from the 'Downloads' section of the Blender Website. Here `2.##` is the Blender version, `#.#.#` is the glibc version installed on your computer and `ARCH` is the machine architecture--either `i386` or `powerpc`. Pick the one matching your system, keeping in mind the difference between static and dynamic builds.

Unpack the archive to a location of your choice. This will create a directory named `blender-2.##-linux-glibc#.#.#-ARCH`, in which you will find the `blender` binary.

To start Blender, open a shell and execute `./blender`, of course while running X.

In-depth Instructions

Download the file `blender-2.##-linux-glibc#.#.#-ARCH.tar.gz` from the 'Downloads' section of the Blender Website. Choose to download it (if prompted), select a location, and click "Save". Then navigate to the location you wish Blender to install to (e.g. `/usr/local/`) and unpack the archive (with `tar -xzf /path/to/blender-2.##-linux-glibc#.#.#-ARCH.tar.gz`). If you like, you can rename the resulting directory from `blender-2.##-linux-glibc#.#.#-ARCH` to something shorter, e.g. just `blender`.

Blender is now installed and can be started on the command line by entering `/path/to/blender` followed by pressing the enter key in a shell. If you are using KDE or Gnome you can start Blender using your file manager of choice by navigating to the blender executable and double-clicking on it.

If you are using the Sawfish window manager, you might want to add a line like `("Blender" (system "blender &"))` to your `.sawfish/rc` file.

How to add Blender's icons in KDE

1. Select the Menu Editor from the System sub-menu of the K menu.
2. Select the sub-menu labeled Graphics in the menu list.
3. Click the New Item button. A dialog box will appear that prompts you to create a name. Create and type in a suitable name and click OK.
4. You will be returned to the menu list, and the Graphics sub-menu will expand, with your new entry highlighted. In the right section, make sure the following fields are filled in: Name, Comment, Command, Type and Work Path.
 - The Name field should already be filled in, but you can change it here at any time.
 - Fill the Comment field. This is where you define the tag that appears when you roll over the icon.

- Click the folder icon at the end of the Command field to browse to the Blender program icon. Select the program icon and click OK to return to the Menu Editor.
- The type should be Application.
- The work path should be the same as the Command, with the program name left off. For example, if the command field reads `/home/user/blender-#.##-linux-glibc#.#.#-ARCH/blender`, the Work Path would be `/home/user/blender-#.##-linux-glibc#.#.#-ARCH/`.

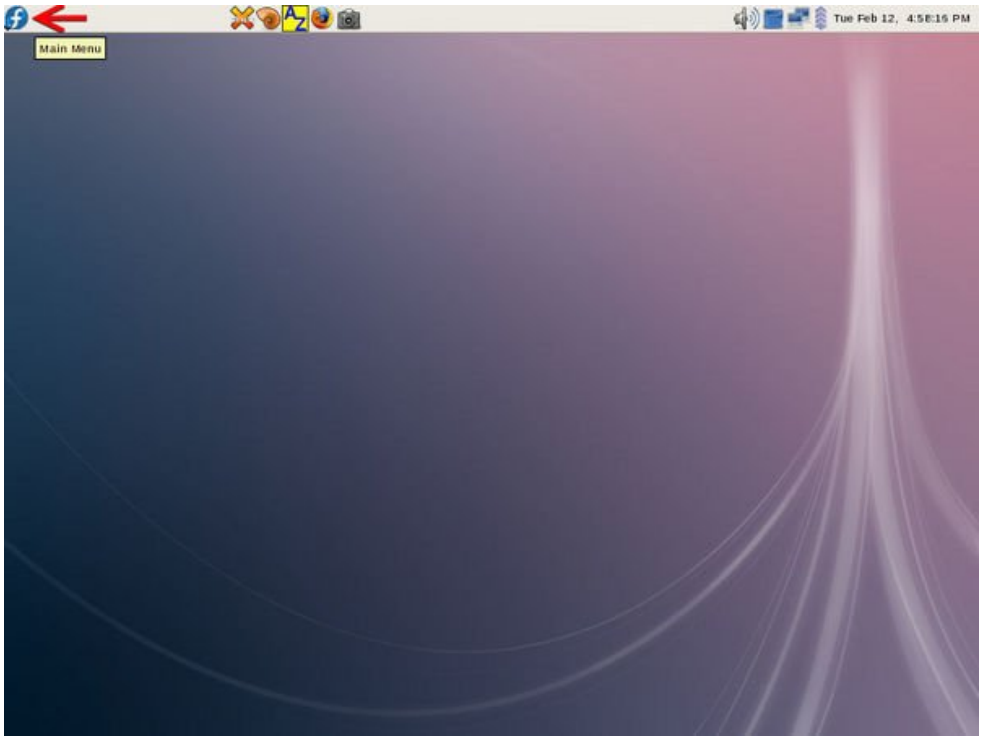
5. Click Apply and close the Menu Editor.

To add a link to Blender on the KPanel, RMB click on a blank spot on the KPanel, then hover over Add. Click Button, then Graphics, and select Blender (or whatever you named the menu item in step 3). Alternately, click on the Configure Panel sub-menu in the K menu, click Add, Button, Graphics, and then Blender.

To add a Desktop icon for Blender, open Konquerer (found on the Panel by default, or in the System sub-menu of the K menu) and navigate to the Blender program icon where you first unzipped it. Click and hold the program icon, and drag it from Konquerer to a blank spot on your Desktop. You will be prompted to Copy Here, Move Here or Link Here; choose Link Here.

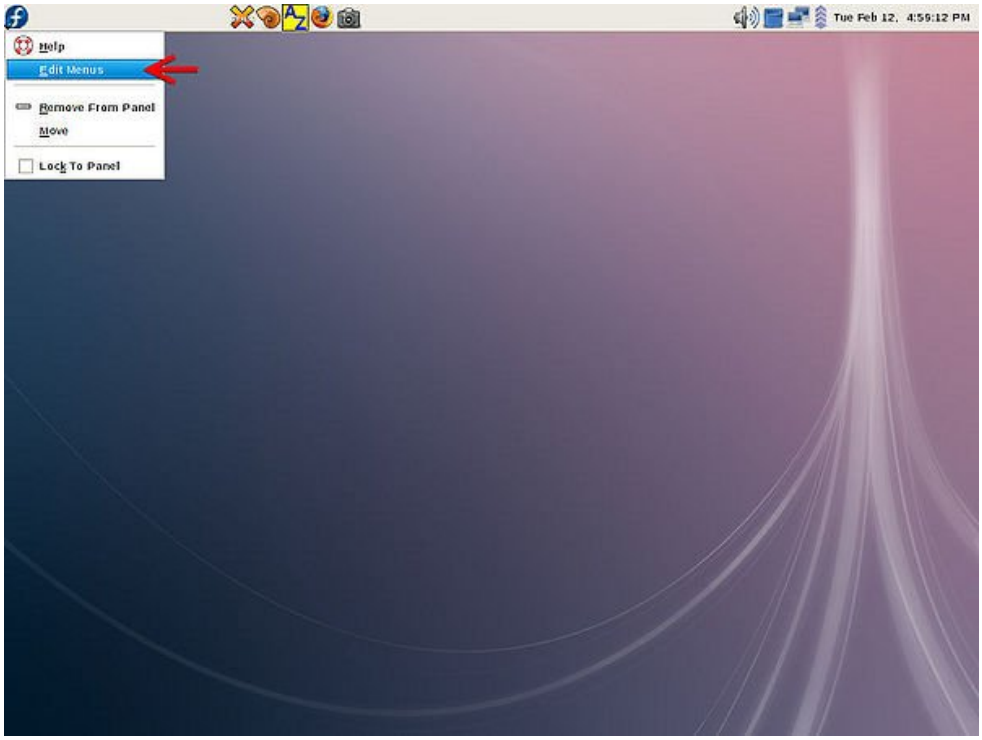
How to add Blender's icons in GNOME

- RMB click the Gnome Main Menu panel (depending on the chosen theme the Icon for the Gnome Main Menu panel could be displayed differently)



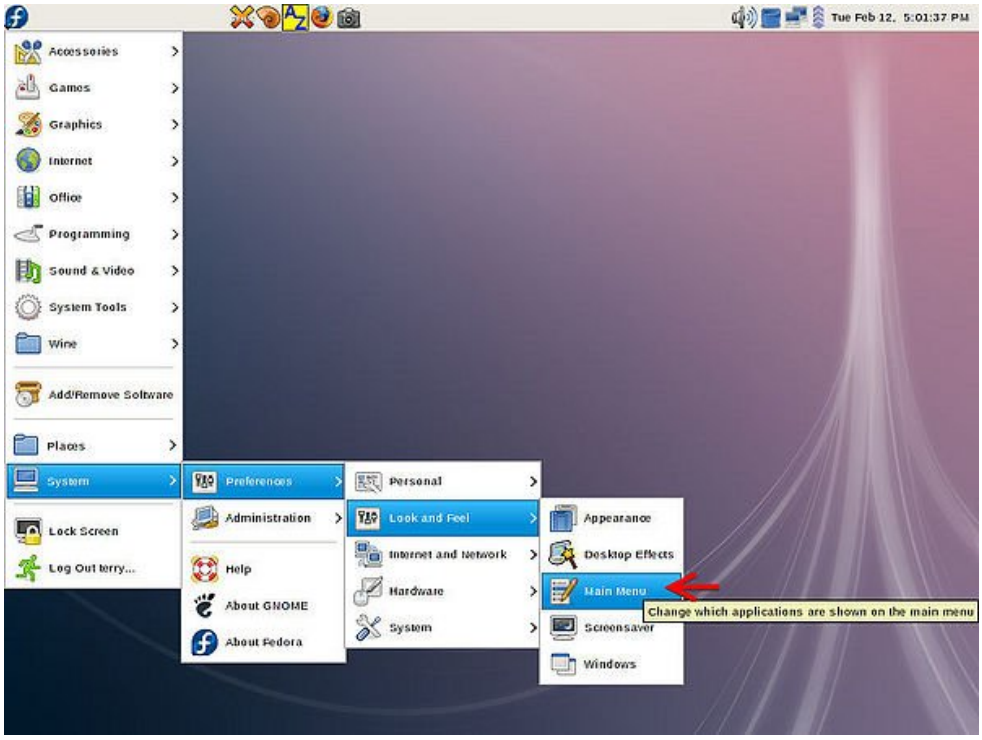
Location of Gnome Main Menu panel

and select Edit Menus from the menu of options that appear,



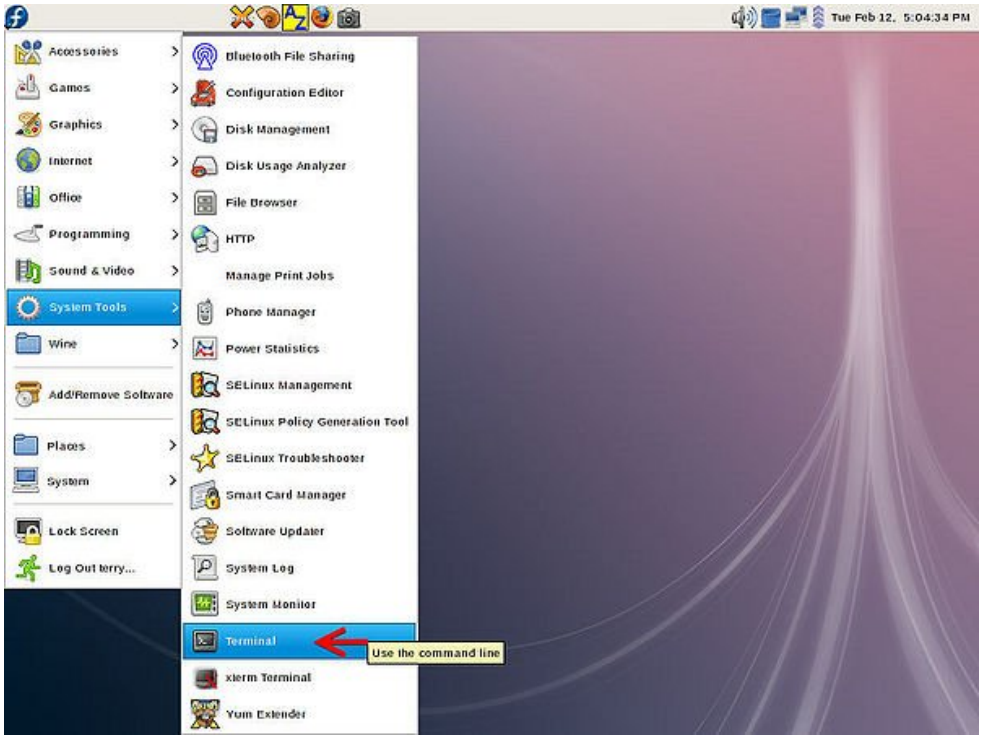
Right clicked panel menu options

or LMB click on the Gnome Main Menu panel and navigate to System > Preferences > Look and Feel > Main Menu (Your menu layout may be different, if so the next option may help).



Location of Main Menu editor

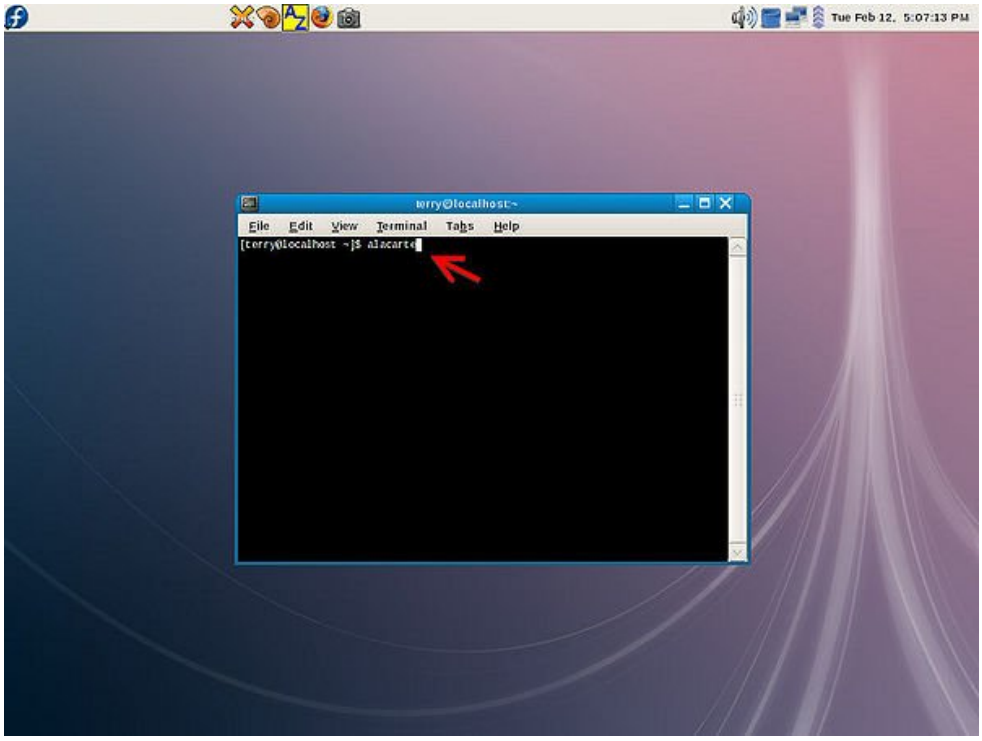
- Yet another method for accessing the Gnome Main Menu editor is to open a Terminal/Console/xterm window



Location of Gnome Terminal

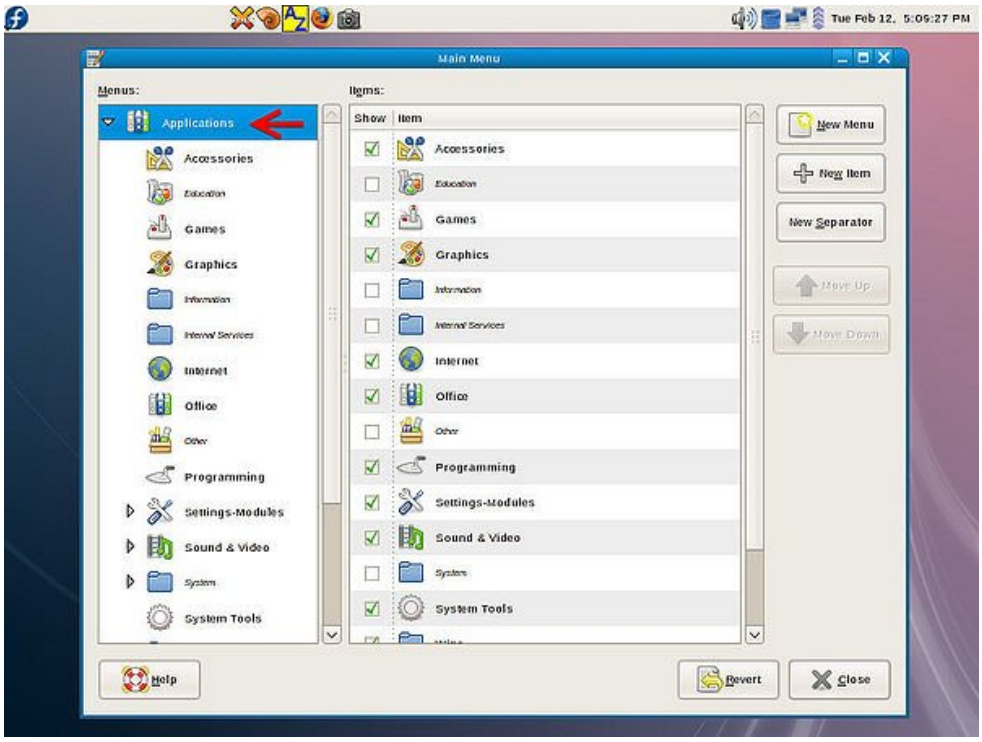
and type the following:

- `alacarte` <PRESS THE ENTER KEY>



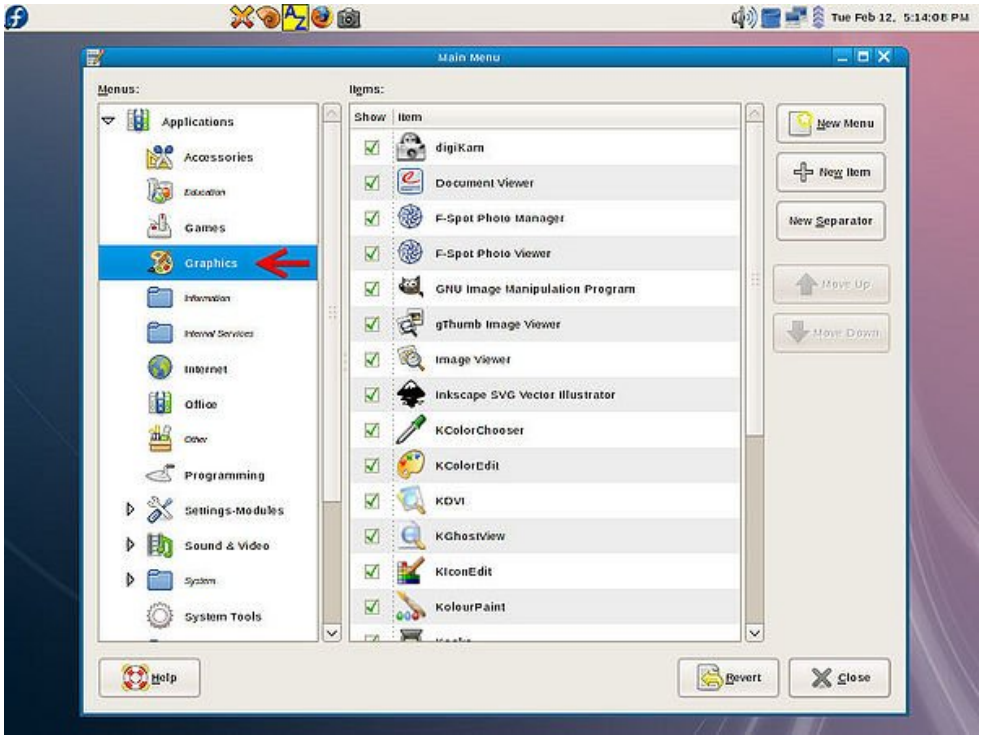
An open Gnome Terminal window

- After using one of the above methods (hopefully) the Main Menu editor is displayed.



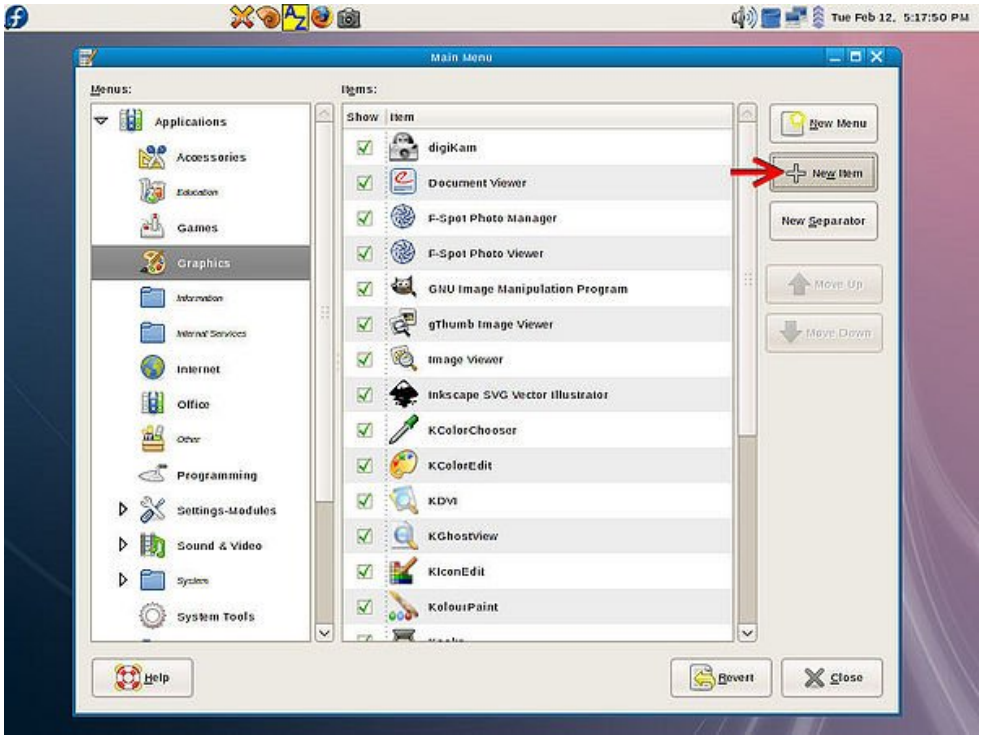
The Main Menu editor windows (alacarte)

- Select the Graphics sub-menu from the Main Menu dialog box (or which ever section you want the Blender icon to be contained in),



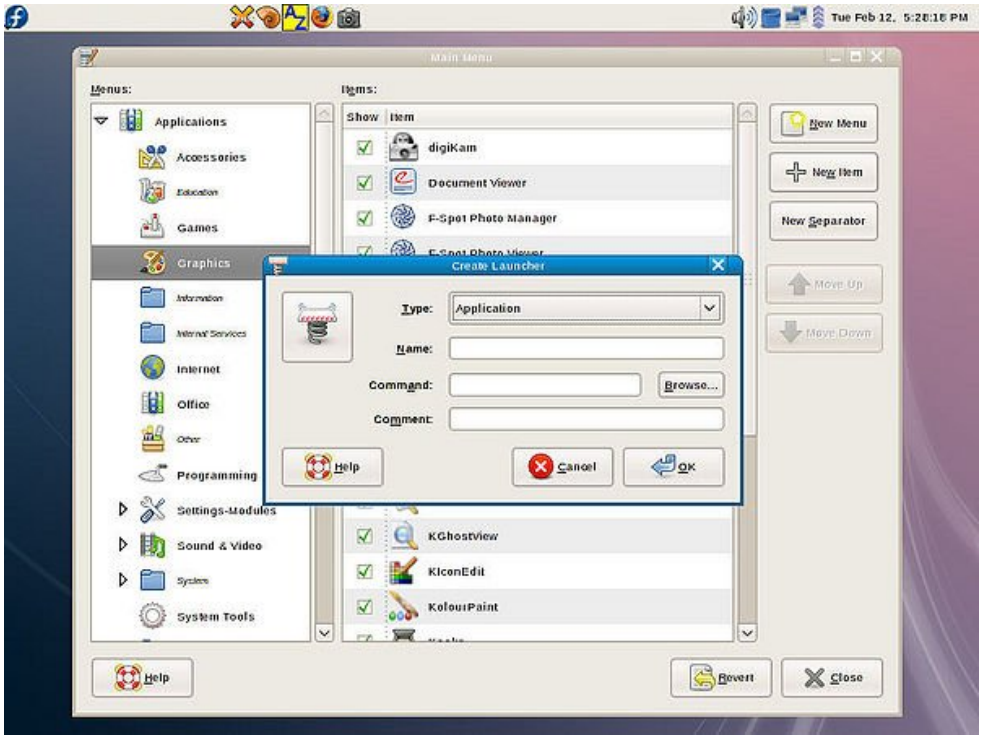
The Main Menu editor windows (alacarte), with the left hand Graphics menu section selected

then click the New Item button.



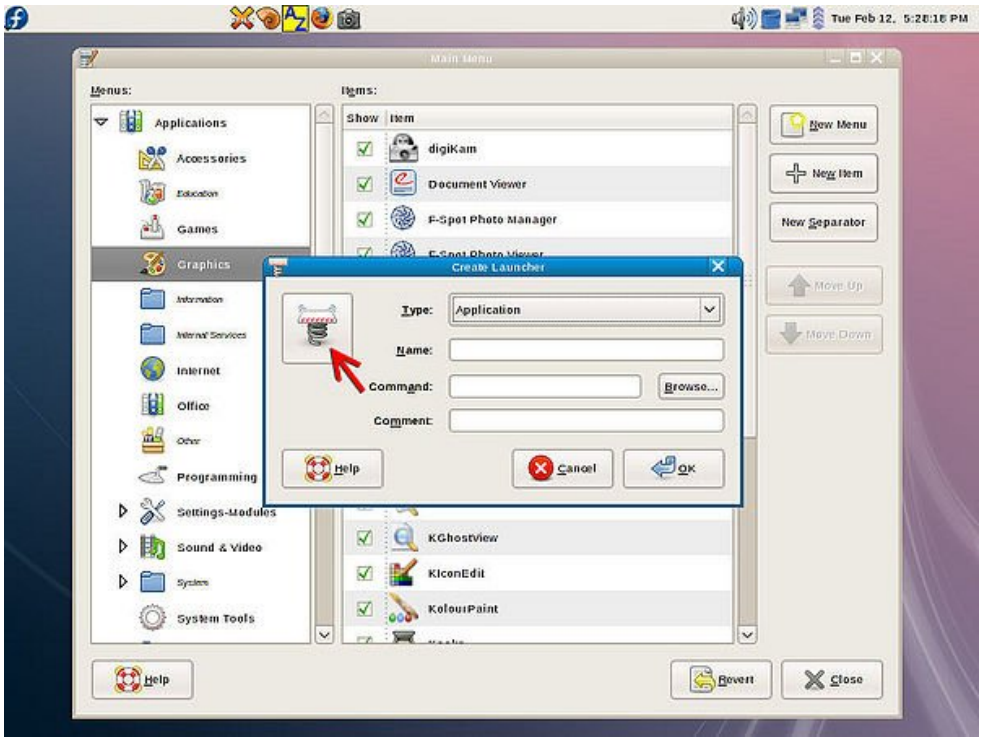
The Main Menu editor windows (alacarte), with the right hand New Item clicked

- In the Create Launcher dialog box, make sure the Type: drop down menu has Application selected.



The Create Launcher dialog box (alacarte)

- In the Create Launcher dialog box also fill in the Name:, Comment: and Command: fields. Fill the Name: field with the program name, for example Blender. You can name this whatever you'd like, this is what appears in the menu, but does not affect the functionality of the program. Fill the Comment: field with a descriptive comment. This is what is shown on the tooltips popups. Fill the Command: field with the full path of the blender program executable/binary, for example, `/home/user/blender-#.##-linux-glibc#.#.#-ARCH/blender`
- Click the icon button to choose an icon (the icon button by default is to the top left within the Create Launcher dialog box and looks like a platform attached to a spring (depending on the chosen theme) or if no icon is selected by the theme the words No Icon may be displayed.



The Icon selection/display box (alacarte)

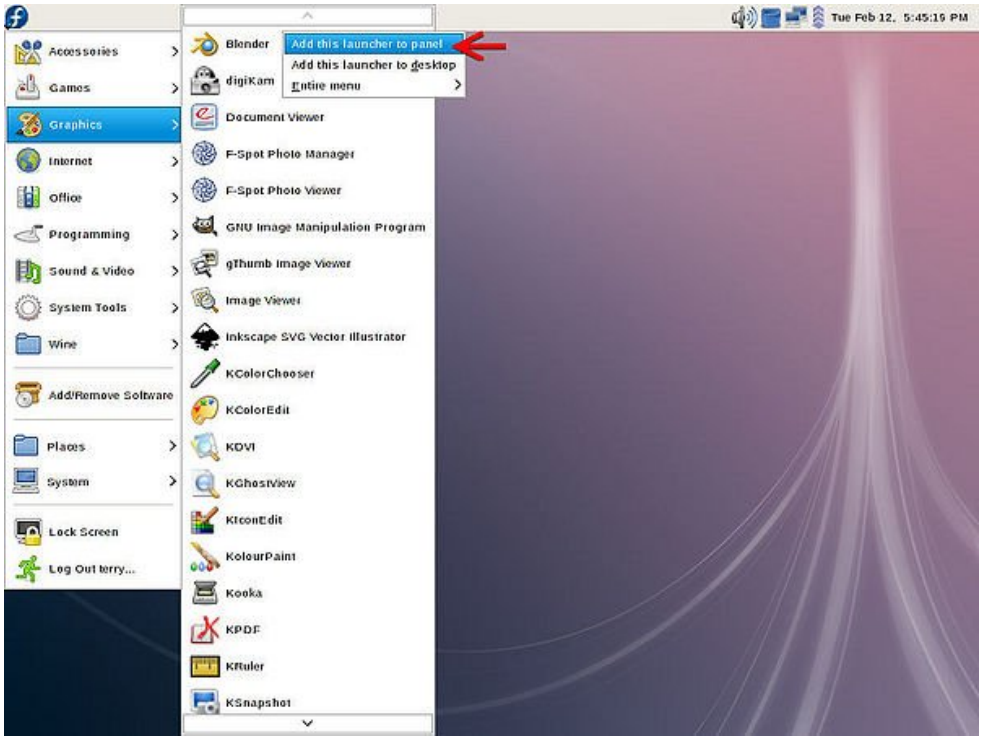
When the mouse is positioned over the icon button it will highlight). There may or may not be an icon for Blender in your default location. You can make one, or look for the icon that goes with KDE. This should be in directory `/opt/kde/share/icons/hicolor/48x48/apps/blender.png` (assuming you installed KDE). If your installation directory is different, you can search for it using this command in a Terminal or Console:

- `find / -name "blender.png" -print` <PRESS THE ENTER KEY>
- If you cannot find your Blender icon then you can use this 2.45 Blender icon for the Gnome Blender icon , just click on the picture and save it to your computer.
- Once you have found the icon you wish to use for Blender, select it in the Browse Icons dialog box and select the Ok button to confirm it.



Browser Dialog displaying available icons in specified location (alacarte)

- Then click the Ok button in the Create Launcher Dialog box to create the new menu and icon item in the Main Menu editor dialog. Make sure the Show item is selected to the left of the newly created Blender entry.
- Click the Close button to close the Main Menu editor.
- Now you should have access to Blender from the Gnome Menu as well as an icon assigned.
- To add a Panel icon for Blender, LMB click on the Gnome Main Menu panel and navigate to the Blender menu entry location in the menu, then RMB click the Blender menu entry and select Add this launcher to panel. Once that is done the Blender icon should appear on the panel.



Right click menu from Main Menu Blender Icon (alacarte)

- To add a Desktop icon for Blender, it is almost the same as adding a Panel icon for Blender but instead of Selecting Add this launcher to panel you instead select Add this launcher to desktop.

OSX

Download the file `blender-2.##-OSX-10.3-py2.#-[platform].zip` from the downloads section of the Blender Website.

- 2.## is the Blender version,
- 10.3 is the minimum OSX version,
- py2.# is the python version, which should be 2.3 for most users, and
- [platform] is either powerpc or i386 (Intel), depending on which type of processor your Mac has.

Python 2.3 is included with the installation. If you wish to use the latest version of Python, please refer to the Python section on this page.

Extract the download by double-clicking the file. This will open a directory with several files.

Mac Users: Since Blender uses OpenGL to render the user interface and Mac OSX draws the entire Desktop with OpenGL as well, you will need to verify that you have sufficient VRAM in your system. Blender will not run with less than 8MB of VRAM. With up to 16 MB VRAM, you will need to set your display to "1000s of colors" (System Preferences -> Displays).

You now can use Blender by double clicking the Blender icon, or drag the Blender icon to the Dock to add its icon there. Blender starts by default in a small window. Hints and tips about the OSX version can be found in the file `OSX tips.rtf` in the installation directory. If Blender doesn't launch, make sure that you downloaded the correct version; oftentimes, newcomers to Blender will accidentally download the Python 2.5 version by accident; try the Python 2.3 version if Blender doesn't seem to launch.

Other supported operating systems

FreeBSD

Download the file `blender-2.##-freebsd-#.#-i386.tar.gz` from the downloads section of the Blender Website. Here `2.##` is the Blender version currently 2.45, `#.#` is FreeBSD version and `i386` is the machine architecture.

To start blender just open a shell and execute `./blender`, of course while running X.

Irix

Download the file `blender-2.##-irix-6.5-mips.tar.gz` from the 'Downloads' section of the Blender Website. Here `2.##` is the Blender version (currently 2.45), Python version 2.4, `6.5` is the Irix version and `mips` is the machine architecture.

To start Blender just open a shell and execute `./blender`, of course while running X. Blender was originally developed for the IRIX platform, but is currently not actively being maintained for all IRIX workstation versions. For some workstations performance troubles have been reported.

Solaris

Download the file `blender-2.##-solaris-2.8-sparc.tar.gz` from the 'Downloads' section of the Blender Website. Here `2.##` is the Blender version (currently 2.45), Python version 2.5, `2.8` is the Solaris version and `sparc` is the machine architecture.

Currently no further instructions for Sun Solaris are available. Please use the Blender Website forums for support.

Other unsupported operating systems

MorphOS

Port on this platform is maintained by Guillaume Roguez. More information on his wiki here: http://www.yomgui.fr/yiki/doku.php/en:dev:blender:start

(French version here: http://www.yomgui.fr/yiki/doku.php/dev:blender:start)

Current status: 2.46 was the latest released version, but not supported anymore due to the full time work on 2.50 release.

Currently no further instructions for MorphOS are available. Please use the Blender Website or MorphOS forums for support.

Configure your Blender

The generic installation of Blender has tons of features and looks pretty cool, too. When you install an upgrade, there are a few things you want to do:

- 'Point' Blender to resources on your machine
- Copy and regression test custom python scripts
- Tell Blender where sequence and texturing plugins are
- Customize your animation, modelling, material, sequence and scripting desktops
- Define your default animation output directory

The top window contains all the User Preferences - click here for more info, including a File Paths tab that you should set up. Then, go into your \Blender\Scripts\ folder and copy the non-distributed scripts into your .blender\scripts directory. Your texture and sequence plugin pathspecs are under User Preferences, and I keep mine in \Blender\bin folders of their own. Your different desktops are selected from the left drop-down menu at the top of the screen. You can size and reconfigure each of these to suit your particular preference (for newbies, the defaults are just fine). If you click the Render button, to top Output directory is where your animations are put (by default), and you might want to point that to your temp directory. Finally, save all your changes with CtrlU.



Info:

The key combination CtrlU saves all the settings of the currently open Blender file into the default Blender file (which is usually called .B.blend). The settings in the default Blender file are read when Blender is first started or when CtrlX is pressed to start a New file. If you accidentally change the settings in your default Blender file there are a few ways of getting back factory default settings:

- Goto the File menu and select Load Factory Settings, once that is done press key combination CtrlU to save the newly loaded factory settings to the Blender default file. If you have an older version of Blender this method may not be available in that case try the second method.

- Delete the .B.blend file (location can vary between operating systems, check your system) and when Blender is restarted it will recreate it using inbuilt default settings.

Setting up a user directory structure

If you are new to setting up Blender on your PC, you may want to stay organized, as you will quickly accumulate many models, textures, pictures, `.blend` files, `.zip` files, scripts, etc. Mushing them all together in one directory leads to confusion, so it is recommended that you spend a few moments creating a few folders to keep stuff organized. The following is a recommendation based on a few years' experience, but is no means the only way to organize your library. There are also free tools to help you manage larger projects (i.e. CVS/Subversion and Verse), and also asset management systems, but those are beyond the scope of this document.

For casual users, a suggested structure to create on your workstation's hard drive is:

- C:\Blender - a shared folder containing the following subfolders:
 - \bin\ - downloaded binaries (installation exe's) and utilities and add-ons such as Yafray, Python, Gocubic, Panocube, Virtual Dub, etc.
 - \examples\ - work done by others (pictures, movies) for offline study
 - \lib\ - a library of reference material (more on this later)
 - \man\ - User manuals, pdf guides such as Blender Basics, videos from experts, quick reference cards and 'how-to' notes you've made
 - \play\ - your own playground; a directory to save .blend files you're just playing around with
 - \script\ - python scripts that are not distributed with Blender
 - \tmp\ - a temporary place for temporary output; a swap space
 - \tut\ - "how to" tutorials collected from the web. There are many videos and web pages out there (save as a complete web page).
 - \util\ - Blender utilities, such as Make Human, World Forge, and Tree Generator.
 - \work\ - And last but not least, if you actually latch onto a meaningful project that maybe evolves out of the playground, put it here.

Folder explanations

The main Folder is `/Blender/`, which I keep on XP under `/Shared Downloads/`. Create a subfolder `/Blender/bin/` to hold the downloaded binaries or `.exe` install file, as well as any other executable programs associated with Blender, such as YafRay, and some nifty DLL's you will run across for extending Blender functionality.

- Library: I know you want to create the world, but there are already a bunch of models and stuff out there on the www that other creative people have created. To hold this wealth of pre-built knowledge, create a library (`/Blender/lib/`) to hold this stuff. Subdirectories under that could be `/mesh` (to hold blend files of meshes), `/tex` to hold texture images, and `/pic` to hold pictures, such as reference pictures. My `/blender/lib/mesh` folder has subfolders `/animal`, `/human`, `/machine`, and `/house`, to name a few, holding blend files that contain models of those types of things. The `/tex` folder has a similar set of folders containing jpg's and even blend files that contain common material settings that are used to color and paint objects. My `/tex` folder contains `/nature`, `/buildings`, `/painted`, and `/metal` subdirectories. The `/pic` folder contains reference pictures of people (Angelina Jolie), faces (my daughter), furniture, my car (a Dodge Viper), and other reference images and concept art that I want to use as reference when modeling.
- Manual and User Guides: Create a `/Blender/man/` file to hold user manuals and guide files in either html, word (`.doc`) and/or `.pdf` formats. There are a few of these floating around. Also, use this folder to save local copies of these wiki pages for off-line reference.
- Tutorials: There are lots of tutorials around and available for downloading. Create a `/Blender/tut/` directory to hold neat tutorials that you find. Some tutorials are hosted by individuals and may disappear, so if you find a tutorial that helps you, download it into this directory.

- Python scripts: Blender uses a scripting language, Python, to extend its functionality. There are dozens of these scripts that can be loaded by Blender. As you find them, save them in a `/Blender/script/` directory, as well as any batch files you write for making backups, etc.
- Utilities: Blender has evolved to the point where there are complete programs that create wondrous things. Keep your Make Human and World Forge utilities (for example) in /Blender/util/.
- Just Do It!: So now YOU need some of your own space, my young padawan. Create `/Blender/play/` and `/Blender/work/` directories to hold play files, and, for when you actually have a meaningful project to work on, a work file. I have used Blender to create a commercial, a documentary on Niger, and a patent (#6,796,205), so I have a subdir under `/Blender/work` for each of those projects. The `/Blender/work/` folder has a folder for each project, and, below that, a set of `/tex`, `/pic`, `/render`, and `/wav` folders to hold textures, pictures, render output, and sound files, respectively. The actual blend files are kept in the `/work/xxx/` folder, where `xxx` is the short name of the project. The `/Blender/play/` folder is loosely organized into `Yafray`, `anim` (animation), `Lighting`, and other folders; basically a trash heap that I rummage around in when I remember that I did something like <insert current problem> but can't remember how to do it again.

Internet Access

Blender is a web-enabled application, in that it has two executables that may access the Internet which your firewall may detect:

- Blender Foundation\Blender\blender.exe - the core application
- Blender Foundation\Python\pythonw.exe - the python script interpreter

In addition, Blender may spawn your web browser (such as Firefox) to access the Internet, for example when you access Help->Manual.

Internet access is not required for normal operation of Blender. However, in order to access the User Manual and other on-line resources, you will need an Internet connection.

Introduction

The user interface enables interaction between the user and the program. The user communicates with the program via the keyboard and mouse, while the program response is on the screen.

If you are new to Blender, you should get a grip of the user interface before you start modeling. The concepts behind Blender's interface are specifically designed for a graphics modeling application and the vast array of features are different from and organized differently than other 3D software packages. In particular, Windows users may need to get accustomed to an alternate manner that Blender handles controls such as button choices and mouse movements.

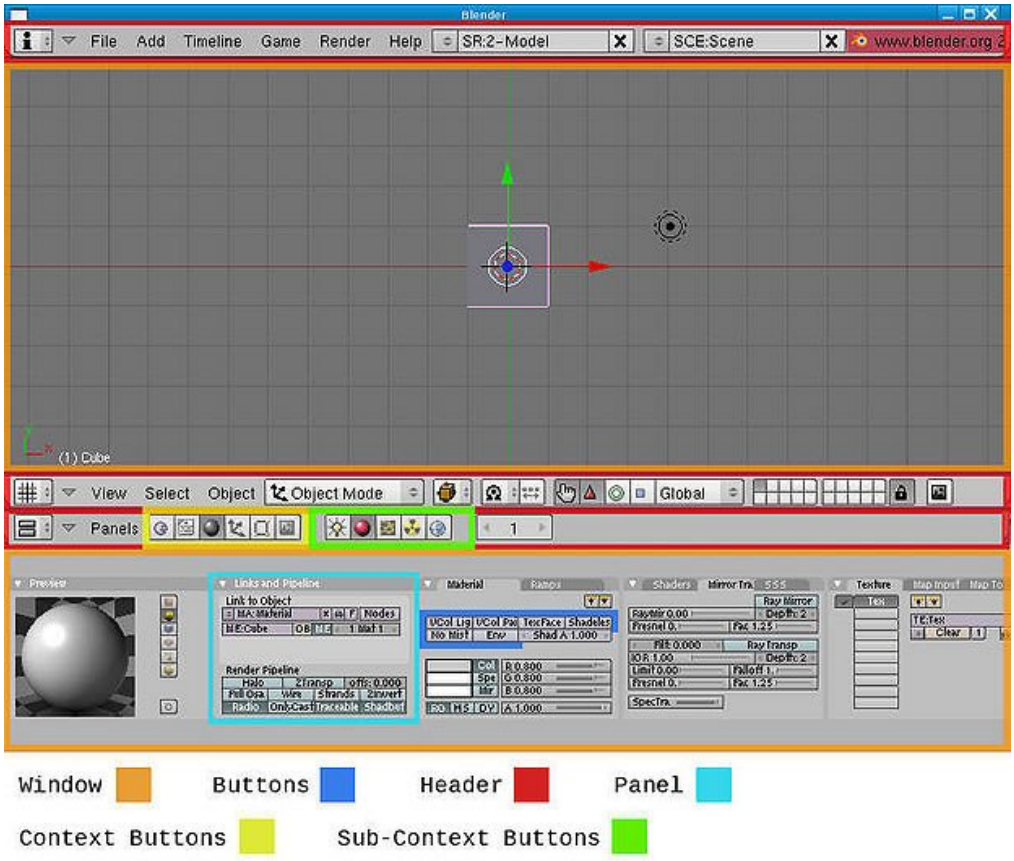
Blender's Interface Concept

This difference is one of Blender's great strengths. Once you understand the Blender methodology, you will find that you can work exceedingly quickly and productively. Some features are familiar, like the top menu bar of "File", "Add"..."Help". However, many other features are quite unheard of in most (if not all) other applications. For example:

- Blender windows cannot overlap and hide each other, one exception being a small number of mini-floating panels which are transparent, fold-able, small, and dock-able.
- Blender relies heavily on keyboard shortcuts to speed up the work.
- Blender's interface is drawn entirely in OpenGL and every window can be panned, zoomed in/out, and its content moved around.
- Your screen can be organized exactly to your taste for each specialized task and this organization can be named and memorized.

These key differences (and many others) make Blender a unique, powerful, and very nimble application, once you take the time to understand it.

The interface can be broken down into several key areas: *Windows*, *Contexts*, *Panels*, and *Buttons* (controls). For example, The **Button window** contains *Context* buttons which show different groups of *Panels* and the *Panels* each show groups of *Buttons*. These principal areas are discussed on the following pages.



Blender interface organisation

About Interface Internationalization

Blender is translated in many languages, and you can select the one you prefer in its User Preferences window, Language & Font section (first enable the International Font button, then the chosen language in the middle drop-down list, and finally enable the UI elements you want translated, using the three Tooltips, Buttons and Toolbox buttons).

If you do not activate any button, everything will remain in default language (i.e. english), so it might be a good idea to only enable the Tooltips one, to have these small chunks of help in your native language, while keeping the interface in its default, english version, as this manual is using this version for screenshots, reference, etc. (and all its translations should do the same!).

Keyboard and mouse

This chapter gives an overview of the general mouse and keyboard usage in Blender and the conventions used in this Manual to describe them, as well as tips on how to use non-standard devices.

Conventions in this Manual

This manual uses the following conventions to describe user input:

- The mouse buttons are called LMB (left mouse button), MMB (middle mouse button) and RMB (right mouse button).
- If your mouse has a wheel, MMB refers to clicking the wheel as if it were a button, while Wheel means rolling the wheel.
- Hotkey letters are shown in this manual like they appear on a keyboard; for example G which refers to the lowercase "g". When used, the modifier â § Shift is specified just as the other modifier keys, Ctrl and/or Alt ; this gives, for example, CtrlW or â § ShiftAltA.
- 0 NumPad to 9 NumPad, + NumPad and so on refer to the keys on the separate numeric keypad. NumLock should generally be switched on.
- Other keys are referred to by their names, such as Esc, â Tab, F1 to F12.
- Of special note are the arrow keys, ←, → and so on.

General Usage

Blender's interface is designed to be best used with a three-button mouse. A mouse wheel is quite useful, but not essential.

Because Blender makes such extensive use of both mouse and keyboard, a golden rule has evolved among Blender users: **Keep one hand on the mouse and the other on the keyboard**. If you normally use a keyboard that is significantly different from the English keyboard layout, you may want to think about changing to the English or American layout for your work with Blender. The most frequently used keys are grouped so that they can be reached by the left hand in standard position (index finger on F) on the English keyboard layout. This assumes that you use the mouse with your right hand.

Mouse Button Emulation

It is perfectly possible to use Blender with a two-button mouse or an Apple single-button Mouse. The missing buttons can be emulated with key/mousebutton combos. Activate this functionality in the User Preferences, View and Controls Context, Emulate 3 Button Mouse button.

The following table shows the combos used:

	2-button Mouse	**Apple Mouse**
LMB	LMB	LMB (mouse button)
MMB	AltLMB	â ¥ OptLMB (Option/Alt key + mouse button)

RMB	RMB	â CmdLMB (Command/Apple key + mouse button)

All the Mouse/Keyboard combinations mentioned in the Manual can be expressed with the combos shown in the table. For Example, â § ShiftAltRMB becomes â § ShiftAltâ CmdLMB on a single-button mouse.

NumPad Emulation

The Numpad keys are used quite often in Blender and are not the same keys as the regular number keys. If you have a keyboard without a Numpad (e.g. on a laptop), you can tell Blender to treat the standard number keys as Numpad keys in the User Preferences, System & OpenGL Context, Emulate Numpad button. A detailed description can be found on this BsoD page.

The Window System

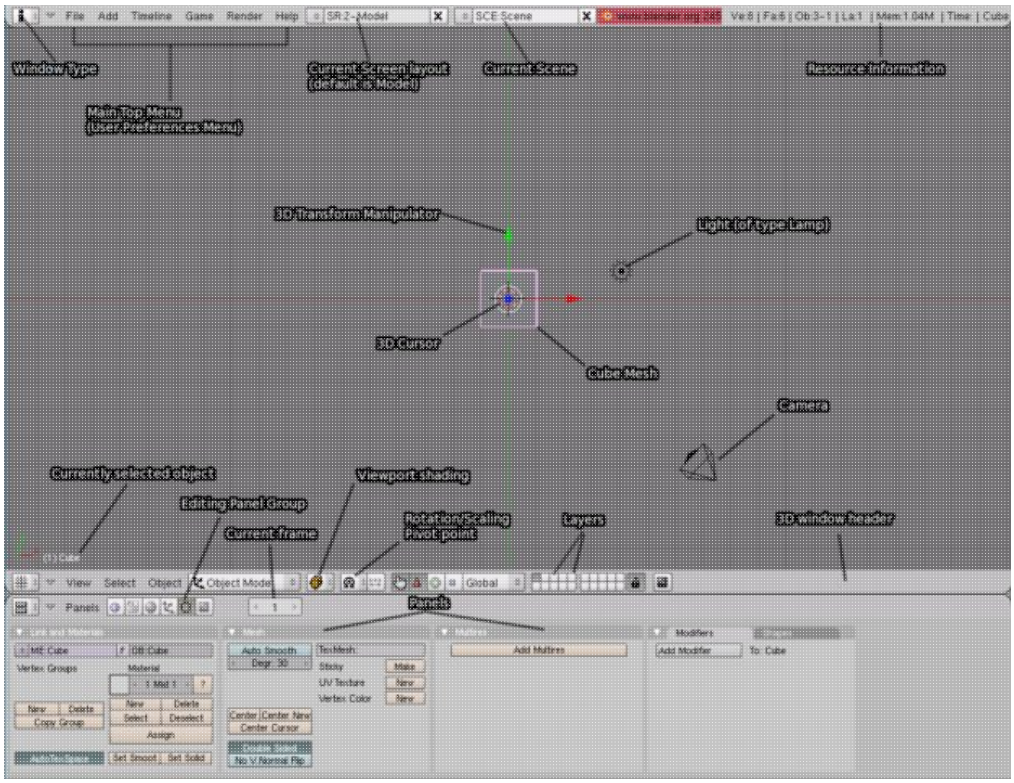
When you start Blender you should see this sequence of events:

- a console window appears
- shortly after, the main user interface window will display the default Blender scene
- you may see a splash screen announcing the Blender version: it usually disappears as soon as you move your mouse so if you move the mouse at the startup too quickly you may not notice the splash screen.

Every window you see can be further broken down into separate areas (see this section): below we describe the default scene.

The default scene

The default Blender scene shows the screen you should get after starting Blender for the first time.



The default Blender scene.
By default it is separated into three windows:

- The main menu at the top is the header part of the User Preferences window
- A large 3D window (3D Viewport window)
- The Buttons Window (at the bottom)

As an introduction we will cover a few of the basic elements.

The main top menu

Check the image above for reference.

- **Window Type**: Allows you to change what kind of window it is. For example, if you want to see the Outliner window you would click and select it.
- **Current Screen (default is Model)**: By default, Blender comes with several pre-configured Screens for you to choose from. If you need custom ones, you can create and name them.
- **Current Scene**: Having multiple scenes present allows for you to break up your work into organized patterns.
- **Resource Information**: Gives you information about application and system resources. It tells you how much memory is being consumed based on the number of vertices, faces and objects in the selected scene. It is a nice visual check to see if you are pushing the limits of your machine.

3D Window View

Check the image above for reference.

- **3D Transform Manipulator**: Is a visual aid in transforming objects. Objects can also be transformed (grabbed/moved - rotated - scaled) using the keyboard shortcuts : (G/R/S); CtrlSpace will display the manipulator pop-up. The manipulator visibility can also be toggled by clicking the "hand" icon on the toolbar. The translation/rotation/scale manipulators can be displayed by clicking each of the three icons to the right of the hand icon. â § ShiftLMB -clicking an icon will add/remove each manipulator's visibility.
- **3D Cursor**: Can have multiple functions. For example, it represents where new objects appear when they are first created; Or it can represent where the base of a rotation will take place.
 Here is the 3D Cursor isolated from the rest of the scene:



- **Cube Mesh**: By default, a new installation of Blender will always start with a Cube Mesh sitting in the center of Global 3D space. After a while, you will most likely want to change the "Default" settings; This is done by configuring Blender as you would want it on startup and then saving it as the "Default" using CtrlU (Save Default Settings).
- **Light (of type Lamp)**: By default, a new installation of Blender will always start with a Light source positioned somewhere close to the center of Global 3D space.
- **Camera**: By default, a new installation of Blender will always start with a Camera positioned somewhere close to the center of Global 3D space and facing it.
- **Currently selected object**: This field shows the name of the currently selected object.

3D Window Header

This is the header for the 3D window. All windows in Blender have a header (and like in this case it may actually be a footer at the bottom of a window but we'll still call it header here anyway). See this section for

an in-depth description about headers.

Check the image above for reference.

- **Viewport shading**: Blender renders the 3D window using OpenGL. You can select the type of interactive shading (called Draw Type: in the Blender shading list) that takes place by clicking this button and selecting from a variety of shading styles. You can select from boxes all the way to complex Textured shading. It is recommended that you have a powerful graphics card if you are going to use the Textured style.

- **Rotation/Scaling Pivot point**: Allows you to select where rotation/scaling will occur. For example, rotation could occur about the object's local origin or about the 3D Cursor's position, amongst many others.

- **Layers**: Make modeling and animating easier. Blender Layers are provided to help distribute your objects into functional regions. For example, one layer may contain a water object and another layer may contain trees, or one layer may contain cameras and lights.

Buttons Window Header

All windows in Blender have a header. See this section for an in-depth description about headers. This is the header for the Buttons window.

Check the image above for reference.

- **Editing Panel Group**: The bottom window displays panels and those panels are grouped. This row of buttons (called Context Buttons) allows you to select which group of panels are shown. Some buttons will display additional buttons (called Sub-Context Buttons) to the right for selection of sub-groups or groups within groups.

- **Current frame**: Blender is a modeling *and* animation application; As such, you can animate things based on the concept of frames. This field shows what the current frame is.

Buttons Window View

- **Panels**: Help group and organize related buttons and controls. Some panels are visible or invisible depending on what type of object is selected.

Changing Window Frames

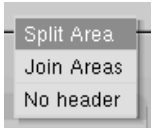
Maximizing a window

You can maximize a window to fill the whole screen with the View → Maximize Window menu entry. To return to normal size, use the View → Tile Window. A quicker way to achieve this is to use â § ShiftSpace, Ctrl↓ or Ctrl↑ to toggle between maximized and framed windows.

Splitting a window

Blender 2.50+
In the upper right hand corner of a window is the window splitter widget, and it looks like a little ridged thumb grip. It both splits and combines window panes. When you hover over it, your cursor will change to a cross. Left-click and drag it to the left to split the pane vertically, or downward to split it horizontally.

Blender 2.49+



The Split menu
You can split a window frame (thus creating a new, smaller window) by *focusing* the window you want to split (moving the mouse to its edge), clicking the vertical or horizontal border with MMB or RMB, and selecting Split Area (*The Split menu for creating new windows.*). You can now set the new border's position by moving your mouse to the desired position, and clicking with LMB; or you can cancel your action by pressing Esc. A click on MMB toggles between vertical and horizontal splitting. The new window will start as a clone of the window you split. It can then be set to a different window type, or to display the scene from a different point of view (in the case of the 3D Window).

Joining two windows

Blender 2.50+
In order to merge two windows, they must be the same dimension in the direction you wish to merge. For example, if you want to combine two windows that are side-by-side, they must be the same height. If the one on the left is not the same as the one on the right, you will not be able to combine them horizontally. This is so that the combined window space results in a rectangle. The same rule holds for joining two windows that are stacked on top of one another; they must both have the same width. If the one above is split vertically, you must first merge those two, and then join the bottom one up to the upper one.

To merge the current window with the one to the right, hover your mouse cursor over the window splitter (that little thumb grip in the upper right-hand corner of the window pane). When your cursor changes to the cross, LMB click and drag to the right to begin the process of combining the window with the one to the right. The window to the right will get a little darker, overlaid with an arrow pointing to the right. This indicates that the left (current) window type will "take over" that darkened window pane space. Click to make it happen. If you want the reverse to occur, move your mouse cursor back into the original (left) window, and the original window will instead get the arrow overlay; click and then the right window will take over the original window.

Similarly for combining with the window above; click the thumb grip and move your mouse straight up into the window above, and click to activate.

Blender 2.49+
You can join two windows into one by clicking a border between two windows with MMB or RMB and choosing Join Areas. Then you'll be prompted to click on one of the two windows; the one you click will disappear, while the other will be expanded to cover the full area of both windows. If you press Esc or RMB before clicking on one of the windows, the operation will be aborted.

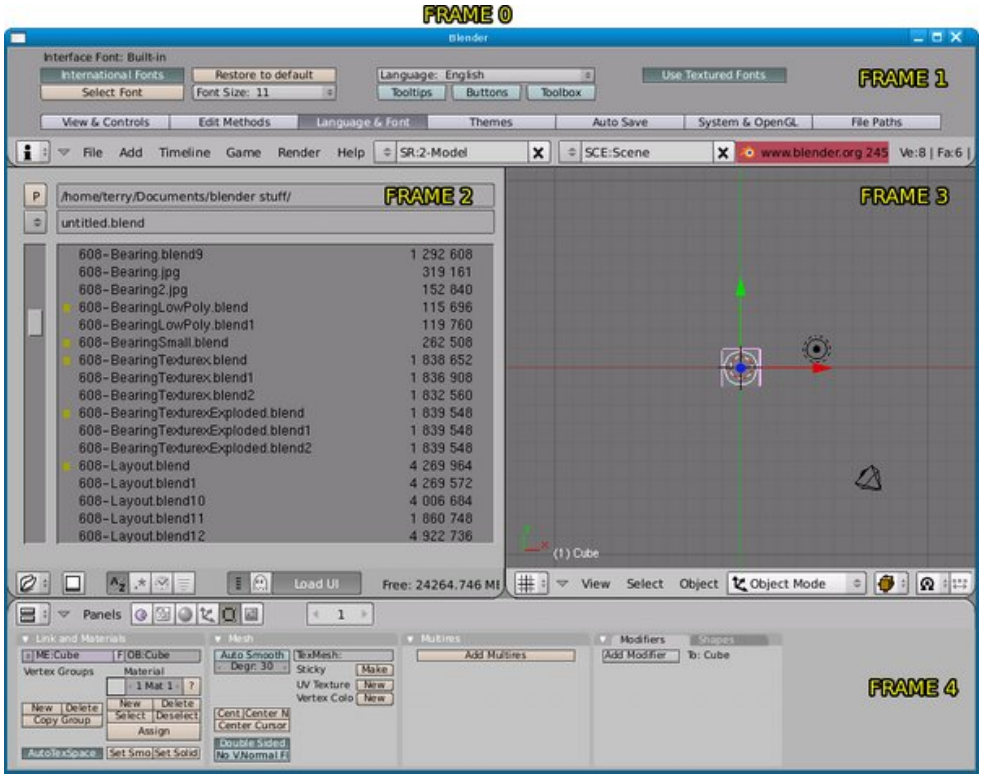
Changing window size

You can resize windows by dragging their borders with LMB. Simply move your mouse cursor until it changes to a double-headed arrow, and then click and drag.

Application Frame

Blender allows the layout of various parts of its interface to be altered in terms of size and position of its window frames; However when using window frame actions such as minimizing and maximizing a window frame, all actions are constrained to the current Application Frame dimensions (also known as the **Top Level Frame**, **Window Manager Frame** or **Frame 0**), which is provided by the operating system and is placed around the Blender application as a whole. For example if you currently have your Application Frame only taking up half of your screen and want it to take up all of your screen you would need to click on the outer Application Frame controls for maximizing windows, rather than using one of the possible Blender key combinations such as Control↑. Using Control↑ while over Frame 2 for example would only make Frame 2 fill the entire space of the Application Frame, not the entire screen (unless the Application Frame was already filling the entire screen). In the screenshot below the Application Frame is indicated by Frame 0 and is light blue with the title Blender in the center of it; Be aware that the Applcation Frame can be different in style, colour and layout and may not be present at all, depending on both the operating system you are running Blender in and the settings used by Blender when it is executed.

Most of the time in this Manual the Application Frame is not shown to both save space and prevent confusion as different operating systems can have different Application Frame layouts.



Blender Application Frame controls.

Interface Items:

Labels in the interface buttons, menu entries, and in general, all text shown on the screen is highlighted in this book like this.

The Window Header

Most windows have a header (the strip with a lighter grey background containing icon buttons). We will also refer to the header as the window *ToolBar*. If present, the header may be at the top (as with the Buttons Window) or the bottom (as with the 3D Window) of a window's area.

Sample Window Headers using the Default Theme



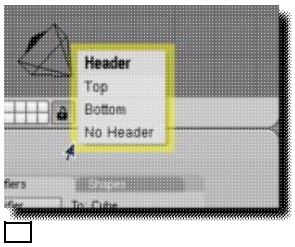
Sample Window Headers using the Rounded Theme



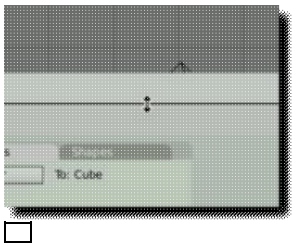
If you move the mouse over a window, its header changes to a lighter shade of grey. This means that it is "focused"; All hotkeys you press will now affect the contents of this window.

Hiding a header

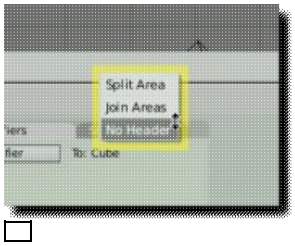
There are a few ways to hide a Window Header from a window:

- 

 Screenshot showing the Header popup menu (highlighted in yellow); The result of RMB clicking the Header Window.
 You can hide a particular window's header by moving your mouse over the Window Header that you wish to hide; Then with the mouse cursor still over the Window Header, click RMB to display a popup menu with the name Header; The Header popup menu has the options, Top, Bottom, No Header, select the No Header menu option to hide the Window Header.

- 

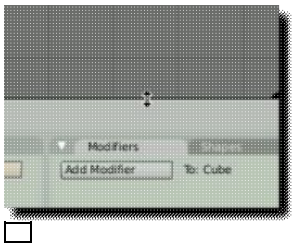
 Mouse cursor positioned over the window frame/border showing the UpDown arrow icon.
 Another method of hiding a particular window's header is to move your mouse over the dividing frame/border next to the Window Header that you wish to hide (which can be just above or just below the Window Header depending on its position), when the mouse cursor is positioned correctly it will display as upward and downward pointing arrows.

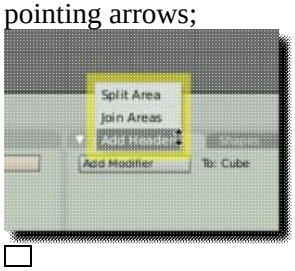


Popup menu that results from RMB clicking on the dividing frame/border. With the No Header menu item selected.
When the upward and downward pointing arrows are displayed RMB click; A popup menu will be displayed with the options Split Area, Join Areas, No Header, select the No Header menu option to hide the Window Header.

Showing a header

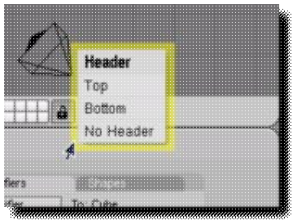
Once a Window Header has been hidden, to redisplay it, do the following:

- 

 Mouse cursor positioned over the Window frame/border of a window with its header removed.
 Move your mouse over the dividing frame/border of the Window Header you wish to unhide (which can be just above or just below the Window Header (that you previously hid) depending on its position), when the mouse cursor is positioned correctly it will display as upward and downward pointing arrows;
- 

 Popup menu that results from clicking RMB on the dividing frame/border. With the Add Header menu item selected.
 When the upward and downward pointing arrows are displayed click RMB or MMB; A popup menu will be displayed with the options Split Area, Join Areas, Add Header, select the Add Header menu option to add the Window Header back to the window.

Header position



Screenshot showing the Header popup menu (highlighted in yellow); The result of RMB clicking the Header Window.
The Window Header can be displayed at the Top or Bottom of the frame. To set a window header's position, RMB click on the Window Header and choose Top or Bottom from the Header popup menu.



Theme colours:

Blender allows for most of its interface colour settings to be changed to suit the needs of the user. If you find that the colours you see on screen do not match those mentioned in the Manual then it could be that your default theme has been altered. Creating a new theme or selecting/altering a pre-existing one can be achieved by selecting the User Preferences window and clicking on the Themes section of the window.



The User Preferences window, Theme section selected.

Window type button

The icon at the left end of a header, with a click of the LMB , allows selection of one of 16 different window types.

Menus and buttons

Most Window Headers, located immediately next to this first "Window Type" Menu button, exhibit a set of menus. Menus allow you to directly access many features and commands. Menus can be hidden and shown via the triangular button next to them.

All Menu entries show the relevant hotkey shortcut, if any.

Menus and buttons change with Window Type and the selected object and mode. They show only actions which can be performed.

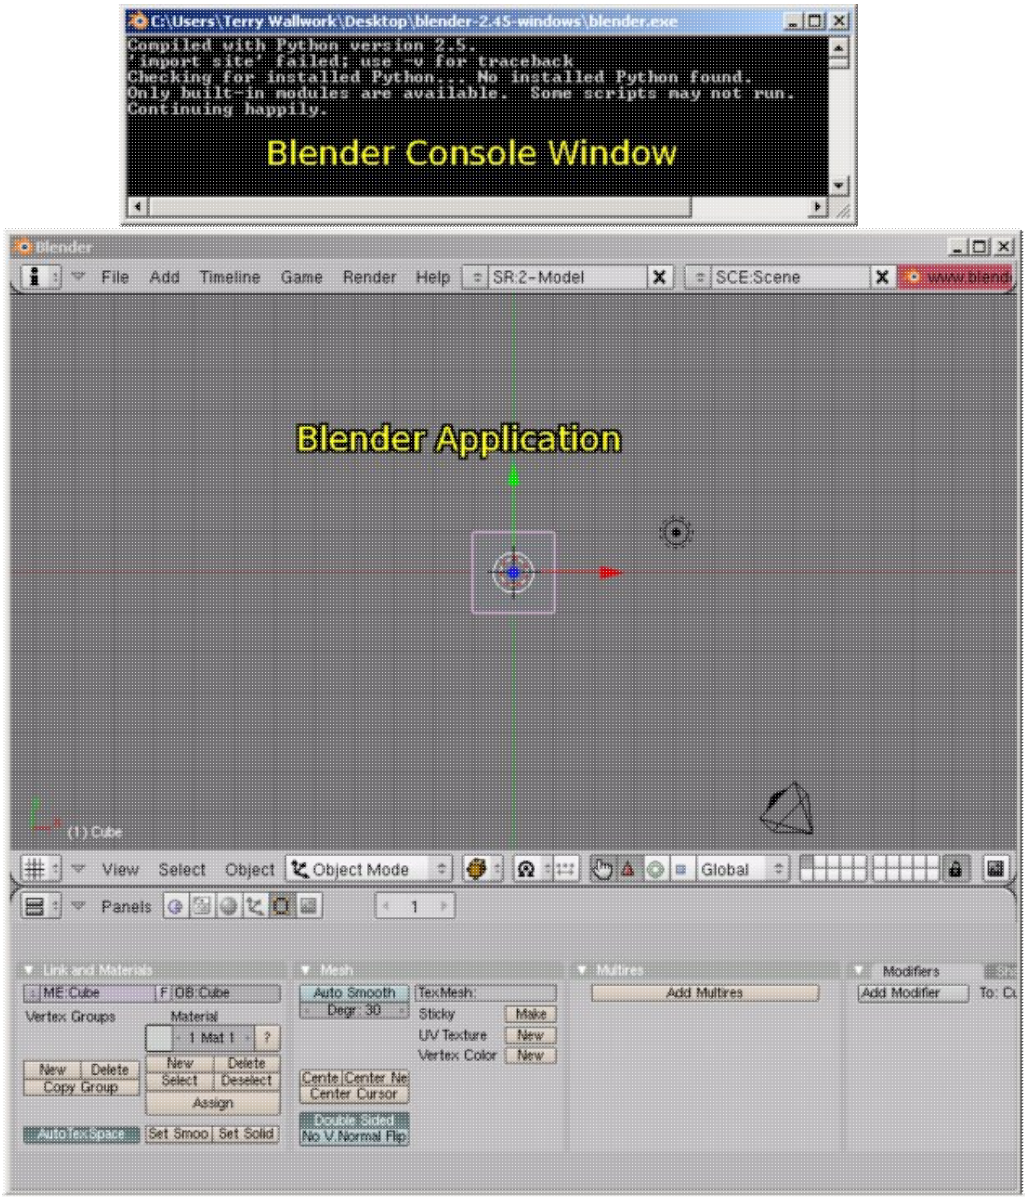
The Console Window

The Console Window is an operating system text window that displays messages about Blender operations, status, and internal errors. If Blender crashes on you, it is a good idea to check the Console Window for clues.

Windows 2000/Xp/Vista

When Blender is started on a Microsoft Windows OS; The Console Window is first created as a separate window on the desktop; Then assuming the right conditions are met, the main Blender Application window should also appear.

This screenshot shows the 2 windows on a Windows Vista OS:



The Blender Console Window and Blender Application.
The Blender Console Window may not be visible, some reasons for this are:

- The Blender Application window may be covering the Console Window. If this is the case just use the Windows task bar to click on the Blender Console Window icon, which should make the Blender Console Window visible.

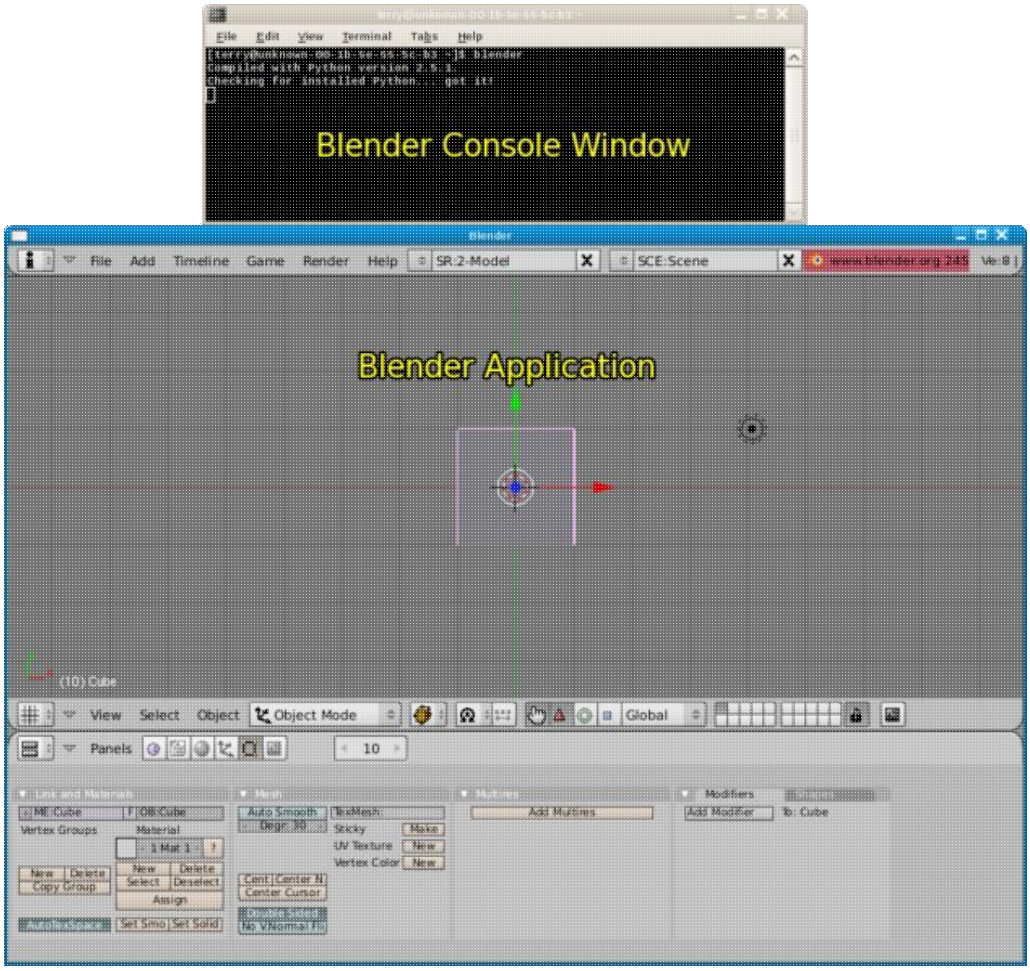
- The Blender Console Window may be minimized/iconifed when Blender starts. If this is the case again, just use the Windows task bar to click on the Blender Console Window icon, which should make the Blender Console Window visible.

Linux

The Blender Console Window in Linux will generally only be visible on the Desktop if Blender is started from a Linux Terminal/Console Window, as Blender uses the Console Window it is started from to display it's Blender Console output.

Most of the different Linux distributions have Blender as one of their applications you can install from their packaging systems. When Blender is installed in this way an icon is usually also installed into their menu systems; Allowing for Blender to be started by clicking an icon rather than having to open a separate Linux Console/Terminal window and start Blender from there; When Blender is started using an icon rather than being started from a Terminal window, the Blender Console Window text will most likely be hidden on the Terminal that XWindows was started from.

This screenshot shows a Linux Terminal/Console Window from which Blender is started; Resulting in Blender outputting it's Console text to it:



Blender in Linux started from a Terminal



Closing the Blender Console Window:

The Blender Console Window must remain open while Blender is executing; If the Blender Console Window is closed then the Blender Application window will also close, and any unsaved Blender work will be lost! The MS DOS command windows and Blender Console Window can look similar, so always make sure that you are closing the correct window (or save your work often in Blender, CtrlW is your friend!)

Console Window Status & Error Messages

The Blender Console Window can display many different types of Status & Error Messages; These can range in level from Trivial (informing the user what Blender is doing, but having no real impact on Blender's ability to function) to Critical (serious errors which will most likely prevent Blender carrying out a particular task and may even make Blender non-responsive/shutdown completely). The Blender Console Window messages can also originate from many different sources (Internally from within the Blender code, Externally from Python scripts which Blender executes, and from varied types of Plugins, to mention a few).

Here is a list of some of the Blender Console Window messages:

- Compiled with Python version X.Y.

 Blender has support for a scripting language called Python; There are many different versions of Python. When Blender software is compiled (programmer's term for building software), it can be compiled to expect a particular version of Python at the version reported on the Blender Console Window. So this message reports the current version of Python the current version of Blender will use when running.

- Checking for installed Python... got it!

 Blender can use the Python language in two different ways, depending on how your system is configured. If you have a fully-fledged version of Python installed on your system, and it is a version that can be used by Blender; Blender will then use the fully-fledged version of the Python interpreter. This allows for more features of Python scripts to be used from within Blender.

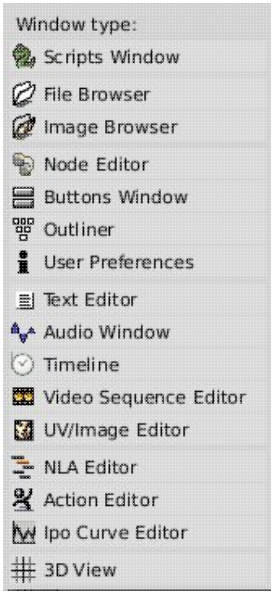
- Checking for installed Python... No installed Python found.

 If Blender cannot find a fully-fledged version of Python on your system or the version it finds is not able to be used; Blender will use an Internal (cut down) version of Python called PyBlender. Even though the Internal version of Python is less feature rich, for the most part it is able to carry out most of the tasks required of it by Blender. If you come across scripts which seem not to work correctly, it may well be that they require a full version of Python to be used successfully; It could also be that the script you're trying to run was written for a different version of Blender/Python. If you wish access to the widest range of Python functionality then there are a few ways to obtain it. One way is to go to http://www.Python.org website and download the version of Python you require. The Windows version of Python comes with a simple-to-use installation program. In Linux you are likely to have Python fully installed already, but if not, you can either compile it and install it manually (often not very easy), or, if you're using a common Linux distribution, have your Linux packaging system install and setup Python for you (usually much easier).

- malloc returns nil()

When Blender carries out tasks that require extra memory (RAM), it calls a function called malloc (short for memory allocate) which tries to allocate a requested amount of memory for Blender. If, however, the amount of memory requested by Blender cannot be satisfied, malloc will return nil/null/0 to indicate that it failed to carry out the request. If this happens Blender will not be able to carry out the tasks required of it by the user. This will most likely result in Blender shutting down or operating very slowly and non-responsively. If you want to avoid running out of memory you can intall more memory in your system, or reduce the amount of detail in your Blender models, or shut down other programs and services which may be taking up memory that Blender could use.

Window types



The window type selection menu.
The Blender interface, the rectangular window provided by your operating system, is divided up into many rectangular window frames. Each window frame may contain different types of information, depending upon the Window type.

Each window frame operates independently of the others, and you can have the same type of window in many frames. For example, you may have several 3D windows open but each looking at the scene from a different perspective. You can split and merge and resize window frames to suit whatever you are working on. You can also arrange some window frames to show with or without a header to save screen space.

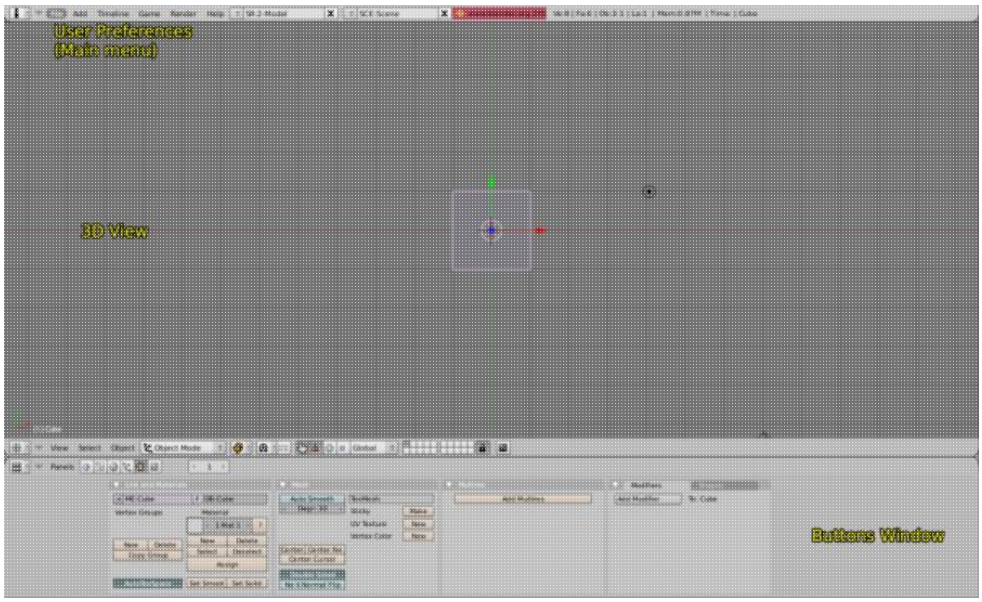
Window types are broken up by functionality:

- Scripts window - user interface for running Python scripts that extend Blender functionality.
- File Browser - for storage and retrieval, especially of .blend files
- Image Browser - search your computer for images, seen as thumbnails.
- Node Editor - process/enhance images and materials.
- Buttons Window - panels that configure objects and set/select options.
- Outliner - Helps you find and organize your objects.
- User Preferences - customize Blender to your work style and computer.
- Text Editor - keep notes and documentation about your project, and write Python scripts.
- Audio Window - see sound files and correlate them to frames
- Timeline - Controls for animation playback.
- Video Sequence Editor - assemble video sequences into a filmstrip.
- UV/Image Editor - editing of the UVmaps ; edit and paint pictures.
- NLA Editor - manage non-linear animation action sequences.
- Action Editor - combine individual actions into action sequences.
- Ipo Curve Editor - manage animation keys and inter/extrapolation of these.
- 3D View - graphical view of your scene.

You can select the Window type by clicking the window's header *leftmost* button. A pop-up menu displays showing the available Window types, see (*The Window type selection menu*).

For further details about each Window type, click on its hyperlink above or visit the reference section III Windows.

The default Blender screen layout is shown below:



Blender default screen layout

Three Window types are provided in Blender's default screen:

3D View

> Provides a graphical view into the scene you are working on. You can view your scene from any angle with a variety of options; see Doc:Manual/3D_interaction/Navigating for details. Having several 3D Viewports on the same screen can be useful if you want to watch your changes from different angles at the same time.

Buttons Window

> Contains most tools for editing objects, surfaces, textures, Lights, and much more. You will need this window constantly if you don't know all hotkeys by heart (memory). You might indeed want more than one of these windows, each with a different set of tools.

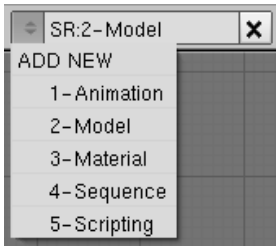
User Preferences *(Main menu)*

> This window is usually hidden, so that only the menu part is visible - see File_operations -> User preferences and Themes for details. It's rarely used though, since it contains global configuration settings that - once saved - don't need to be changed that often. However, the header is frequently used because it provides the only access to a full File menu and to the Add menu.

See also

- Doc:Reference/Buttons/Window

Screens



Layout dropdown
Blender's flexibility with windows lets you create customized working environments for different tasks, such as modeling, animating, and scripting. It is often useful to quickly switch between different environments within the same file. For each Scene, you need to set the stage by modeling the props, dressing them and painting them through materials, etc. In the example picture in Window system, we are in the modeling stage.

To do each of these major creative steps, Blender has a set of pre-defined screens, or window layouts, that show you the types of windows you need to get the job done quickly and efficiently:

1-Animation

Making actors and other objects move about.

2-Model

Creating actors, props, and other objects.

3-Material

Painting and texturing surfaces.

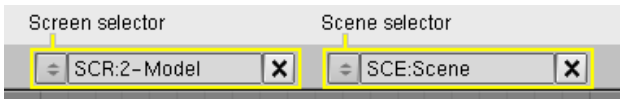
4-Sequence

Editing scenes into a movie.

5-Scripting

Documenting your work, and writing custom scripts.

Blender sorts these screen layouts for you automatically in alphabetical and/or numerical order. The preset screen's names typically start with a number. The list is available via the SCR Menu Buttons in the User Preferences Window header shown in (*Screen and Scene selectors*). To change to the next screen alphabetically press Ctrl→; to change to the previous screen alphabetically, press Ctrl←.



Screen and Scene selectors
By default, each screen layout 'remembers' the last scene it was used on. Selecting a different layout will switch to the layout **and** jump to that scene.

All changes to windows, as described in Window system and Window types, are saved within one screen. If you change your windows in one screen, other screens won't be affected, but the scene you are working on

stays the same in all screens.

Screens Configuration

Adding a new Screen

As you scroll through the Screen list, you will see that one of the options is to **Add New** - namely, add a new window layout. Click () and select ADD NEW. When you click this, a new frame layout is created based on your current layout.

Give the new screen a *name* that starts with a *number* so that you can predictably scroll to it using the arrow keys. You can rename the layout by LMB into the field and typing a new name, or clicking again to position the cursor in the field to edit. For example you could use the name "6-MyScreen". See (*Screen and Scene selectors*).

Deleting a Screen

You can delete a screen by using the Delete datablock button () and confirm by clicking Delete current screen in the pop-up dialog box. See (*Screen and Scene selectors*).

Rearranging a Screen

Use the window controls to move frame borders, split and consolidate windows. When you have a layout that you like, CtrlU to update your User defaults. The buttons window has a special option: If you RMB on its background, you can arrange these panels horizontally (across) or vertically (up and down).

Overriding Defaults

When you save a .blend file, the screen layouts are saved in it. When you open a file, the Load UI button on the file browser header controls whether Blender should use the file's screen layouts, or stick with your current layouts. If Load UI is enabled, the file's screen layouts are used, overriding your defaults.

Additional Layouts

With the dramatic increases in functionality, and as you get better at using Blender, based on what you use Blender for, consider adding some other screen layouts (for a complete workflow):

1-Model

4 3D windows, Buttons window for Editing buttons

2-Lighting

3D windows for moving lights, UV/Image for displaying Render Result, buttons window for rendering and lamp properties and controls.

3-Material

Buttons window for Material settings, 3D window for selecting objects, Outliner, Library script (if used)

4-UV Layout

UV/Image Editor Window, 3D Window for seaming and unwrapping mesh

5-Painting

UV/Image Editor for texture painting image, 3D window for painting directly on object in UV Face Select mode, 3 mini-3D windows down the side that have background reference pictures set to full strength, Buttons window

6-Animation

Ipo Window, 3D Window for posing armature, NLA Window

7-Node

Big Node Editor window for noodles, UV/Image window linked to Render Result

8-Sequence

Ipo Window, VSE window in Image Preview mode, VSE in timeline mode, a Timeline window, and the good old Buttons window.

9-Notes/Scripting

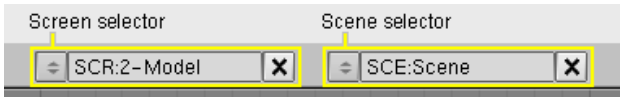
Outliner, Text Editor (Scripts) window

Reuse your Layouts
If you create a new window layout and would like to use it for future .blend files, simply save it as a User default by pressing CtrlU.

Scenes

It is also possible to have several scenes within the same Blender file (to understand what this means, please have a look to the Blender's Library and Data System).

Scenes may use one another's objects or be completely separate from one another. You can select and create scenes with the SCE menu buttons in the User Preferences Window header (*Screen and Scene selectors*).



Screen and Scene selectors

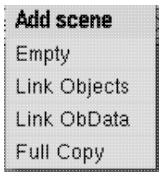
Scenes configuration

Adding a new Scene

You can add a new scene by clicking () and selecting ADD NEW. When you create a new scene, you can choose between four options to control its contents (*Add Scene menu*).

To choose between these very-important options, you need to clearly understand what "Objects," *vs.* "ObData," actually are. "ObData" is geek-speak for "object data." Each Blender graphic element (Mesh, Lamp, Curve, *etc.*) is composed from two parts: an Object and ObData. The Object holds information about the position, rotation and size of a particular element. The ObData holds information that is common to every instance of that particular type of element: *e.g.* meshes, materials lists, and so on. Each Object has a link to its associated ObData, and a single ObData may be shared by many Objects.

The four choices, therefore, determine just how much of this information will be *copied from* the currently selected scene to the new one, and how much will be *shared* ("linked"):



Add Scene menu

- Empty creates an empty scene.
- Link Objects is the shallowest form of copying available. This option creates the new scene with the same contents as the currently selected scene. Nothing is copied: instead, the new scene contains *links to* the Objects in the old scene. Therefore, changes in the new scene will also modify the other, because the Objects used in the new scene *literally are* the same Objects used in the currently selected scene.
- Link ObData creates new, duplicate copies of all of the Objects in the currently selected scene, but each one of those duplicate Objects will have *links to* the same meshes, materials, and so on (the "ObData") as the corresponding Objects in the original scene. This means that you can change the position, orientation and size of the Objects in the new scene without affecting other scenes; but any modifications to the meshes, materials, and so on (the "ObData") will also affect other scenes. (This is because a *single instance of* the "ObData" is now being shared by all of the Objects, in all of the scenes, that link to it. If you want to make changes to an object in the new scene independently of the objects in the other scenes, you'll have to manually make the object in the new scene "single-user" (in the Links and Materials panel of the Editing buttons), which has the effect of making a new independent copy of its ObData.)
- Full Copy is the deepest form of copying available. Nothing is shared. This option creates a fully independent scene with copies of the currently selected scene's contents. Every Object in the original scene is duplicated, and a duplicate, private copy of its ObData is made as well.

Deleting a Scene

You can delete a scene by using the Delete datablock button () and confirm by clicking Delete current scene to the pop dialog box. See (*Screen and Scene selectors*).

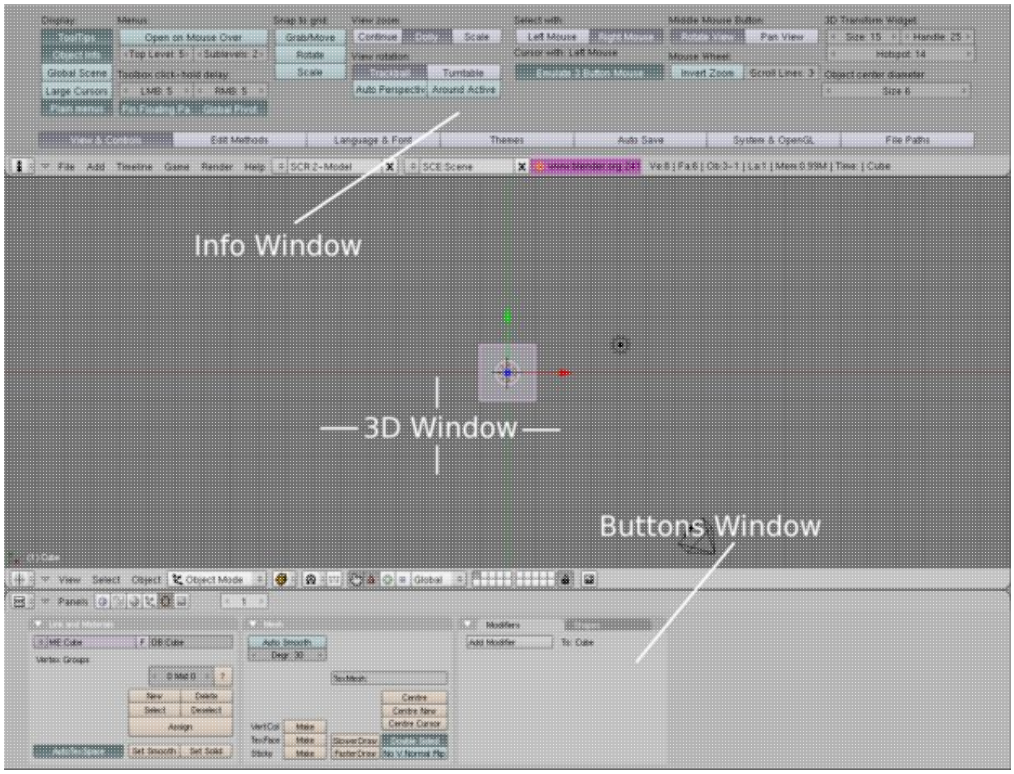
Configuration

The Info window () is where you customize and control Blender. By default this window is located at the top and only the header is visible.



Info window header

To see all of the Info window and its content you need to drag it into view. You can do this by moving the mouse onto the bottom edge of the Info header, or the top of the 3D window, and click the LMB and drag downwards. In picture:Info Visible, the Info window has been made visible at the top.



Info Visible

When viewing all of the Info window you can begin to customize Blender to fit your personality or machine capabilities. For example, you may not like the default theme and switch to the Rounded theme. Or your machine may not be able to handle Vertex Arrays so you switch them off.

For an in depth look at the Info window read the reference section on Info window. There you will find all the details on configuring Blender.

Contexts

The Button Window shows six main Contexts, which can be chosen via the first icon row in the header (*Contexts and Sub-Contexts Example*). Each of these might be subdivided into a variable number of sub-contexts, which can be chosen via the second icon row in the header (*Contexts and Sub-Contexts Example*), or cycled through by pressing the same Context button again:



Contexts and Sub-Contexts Example

- Logic (F4) - Switches to Logic context.
- Script - No shortcut. Switches to Script context.
- Shading (F5) - Switches to Shading context.
 - Lamp - No shortcut.
 - Material - No shortcut.
 - Texture - Shortcut F6.
 - Radiosity - No shortcut.
 - World - Shortcut F8.
- Object (F7) - Switches to Object context.
 - Object - No shortcut.
 - Physics - No shortcut.
- Editing (F9) - Switches to Editing context.
- Scene (F10) - Switches to Scene context.
 - Rendering - No shortcut.
 - Anim/Playback - No shortcut.
 - Sound - No shortcut.

Once the Contexts is selected by the user, the sub-context is usually determined by Blender on the basis of the active Object. For example, with the Shading context, if a Lamp Object is selected then the sub-context shows Lamp Buttons. If a Mesh or other renderable Object is selected, then Material Buttons is the active sub-context, and if a Camera is selected the active sub-context is World.

The Buttons in each context are grouped into Panels.

The menu of available options, shown in a window's header, may change depending on the mode of that window. For example, in the 3D View window, the Object menu in Object mode changes to a Mesh operations menu in Edit mode, and a paint menu in Vertex Paint mode. If you are reading this manual and some menu option is referenced that does not appear on your screen, it may be that you are not in the proper mode, or context, for that option to be valid.

Menus



Top Level
Blender contains many menus each of which is accessible from either the window headers or directly at the mouse's location using HotKeys, see Toolbox.

For example, you can access the Toolbox in the 3D window using either the mouse or the keyboard. From the keyboard you would use the Space. To access it using the mouse just hold down the LMB or RMB buttons for a few seconds and the Toolbox will pop-up. (*Top Level*) is the top most menu of the Toolbox.

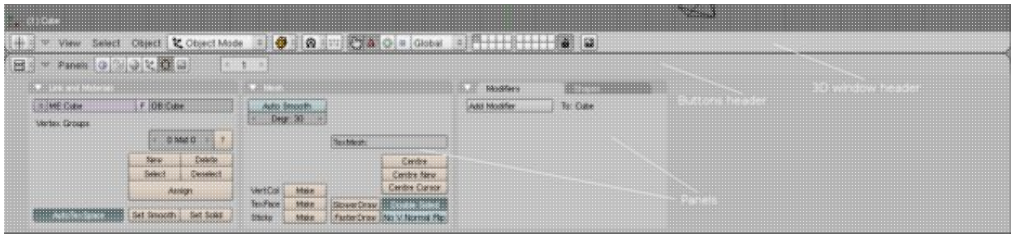
Some menus are context sensitive in that they are only available under certain situations. For example, the Booleans menu is only available in Object Mode using the (W) hotkey. The same hotkey (W) in Edit Mode brings up the Specials menu.

While you are using Blender be aware of what mode and types of object are selected. This helps in knowing what hotkeys work at what times.

If you are using a Mac, because Blender doesn't use the standard OS X menu system, you likely have a redundant menubar at the top. To remove this see this post on Macworld, but beware that it is somewhat complex.

Panels

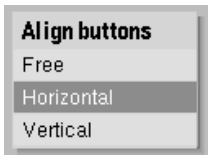
Panels generally appear in the Buttons window and by default the Buttons window is at the bottom; see (*Buttons window*). The Buttons window includes the Button window header and panels.



Buttons window

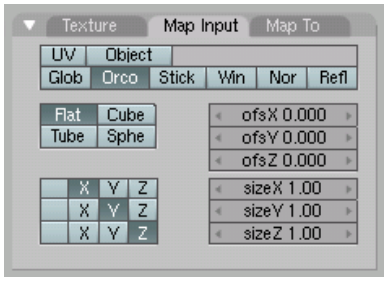
Each button on the Buttons header groups panels together into what is called a Context. And those Contexts are grouped further into Sub-Contexts. For example, all Material panels are grouped under the Shading context and Material sub-context.

The panels are not fixed in position relative to the window. They can be moved around the window by LMB clicking and dragging on the respective panel header.



Button Window Menu.

Panels can be aligned by RMB on the Buttons Window and choosing the desired layout from the Menu which appears (*Button Window Menu.*). Using Wheel scrolls the Panels in their aligned direction and CtrlWheel and CtrlMMB zooms the Panels in and out. Single Panels can be collapsed/expanded by LMB clicking the triangle on the left side of their header.



Panel with Tabs Example.

Particularly complex Panels are organized in *Tabs*. Clicking LMB on a Tab in the Panel header changes the buttons shown in (*Panel with Tabs Example.*). Tabs can be "torn out" of a Panel to form independent panels by clicking LMB on their header and dragging them out. In a similar way separate Panels can be turned into a single Panel with Tabs by dropping one Panel's header into another. **Note** : The colors of the other Panel's borders change (to white by default) when you can release the mouse button to create a tab from the selected Panel. Example

For further details about each panel see the Reference panels section.

Buttons and Controls

Buttons are mostly grouped in the Button Window. But they can appear in other Windows.

Operation Button



An operation button
These are buttons that perform an operation when they are clicked (with LMB, as all buttons). They can be identified by their brownish color in the default Blender scheme (*An operation button*).

Toggle Button



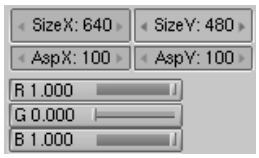
Toggle buttons
Toggle buttons come in various sizes and colours (*Toggle buttons*). The colours green, violet, and grey do not change functionality, they just help the eye to group the buttons and recognize the contents of the interface more quickly. Clicking this type of button does not perform any operation, but only toggles a state.

Some buttons also have a third state that is identified by the text turning yellow (the Emit button in *Toggle buttons*). Usually the third state means "negative," and the normal "on" state means "positive."

Radio Buttons

Radio buttons are particular groups of mutually exclusive Toggle buttons. No more than one Radio Button in a given group can be "on" at one time.

Number Buttons



Number buttons
Number buttons (*Number buttons*) can be identified by their captions, which contain a colon followed by a number. Number buttons are handled in several ways: To increase the value, click LMB on the right of the button, where the small triangle is shown; to decrease it, click on the left of the button, where another triangle is shown.

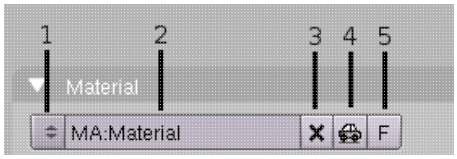
To change the value in a wider range, hold down LMB and drag the mouse to the left or right. If you hold Ctrl while doing this, the value is changed in discrete steps; if you hold â § Shift, you'll have finer control over the values. â µ Enter can be used in place of LMB here.

You can enter a value directly by holding â § Shift and clicking LMB . You can also enter simple equations, like 3*2 instead of 6. Be sure to enter a decimal point if you want a floating point result; otherwise you get an integer (e.g. 3/2 = 1, but 3.0 / 2 = 1.5). Handy geometric constants to remember: pi is 3.14 and the square root of two is 1.414. Press â § Shiftâ Backspace to clear the value; Ctrl← to move the cursor to the beginning; and Ctrl→ to move the cursor to the end. Press Esc to restore the original value. You can copy the value of a button by pointing at it and pressing CtrlC. Similarly you can paste a copied value with CtrlV.

Some number buttons contain a slider rather than just a number with side triangles. The same method of operation applies, except that single LMB clicks must be performed on the left or on the right of the slider, while clicking on the label or the number automatically enters keyboard input mode.

Menu Buttons

Use the Menu buttons to choose from dynamically created lists. Menu buttons are principally used to link DataBlocks to each other. (DataBlocks are structures like Meshes, Objects, Materials, Textures, and so on; by linking a Material to an Object, you assign it.)



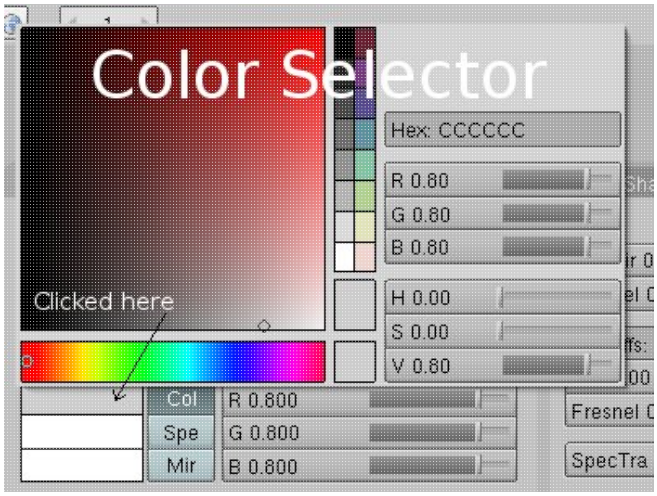
Datablock link buttons
You can see an example for such a block of buttons in (*Datablock link buttons*).

1. The first button (with the tiny up and down pointing triangles) opens a menu that lets you select the DataBlock to link to by holding down LMB and releasing it over the requested item.
2. The second button displays the type and name of the linked DataBlock and lets you edit its name after clicking LMB .
3. The "X" button clears the link.
4. The "car" button generates an automatic name for the DataBlock.
5. And the "F" button specifies whether the DataBlock should be saved in the file even if it is unused (unlinked).

Unlinked objects
Unlinked data is ***not* lost until you quit Blender**. This is a powerful Undo feature. if you delete an object the material assigned to it becomes unlinked, but is still there! You just have to re-link it to another object or press the "F" button.

Color Selector controls

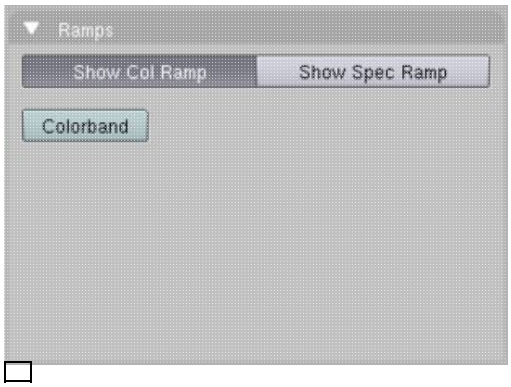
Some controls pop-up a dialog panel. For example, Color controls, when clicked, will pop up a Color Selector dialog; see (*Color Selector*).



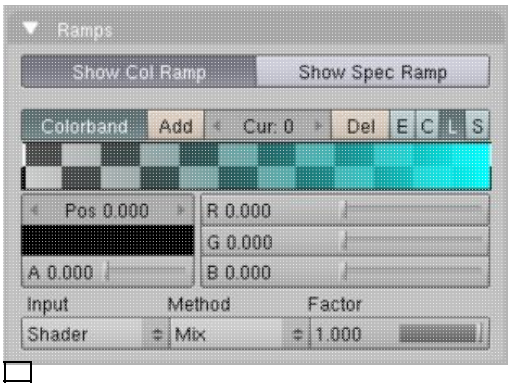
Color Selector

Cascade Buttons

Occasionally, some buttons actually reveal additional buttons. For example, the Ramps panel has a Cascade button called Colorband that reveals additional buttons dealing with colorbanding; see (*Colorband before*) and (*Colorband after*).



Colorband before



Colorband after

Your First Animation in 30 plus 30 Minutes Part I

This chapter will guide you through the animation of a small "Gingerbread Man" character. We will describe each step completely, but we will assume that you have read the interface chapter, and that you understand the conventions used throughout this book.

In Part I of this tutorial we'll build a *still* Gingerbread Man. Then, in Part II, we will make him walk.

Note
For a much more in-depth introduction to Blender that focuses on character animation, check out the

Blender Summer of Documentation Introduction to Character Animation tutorial.





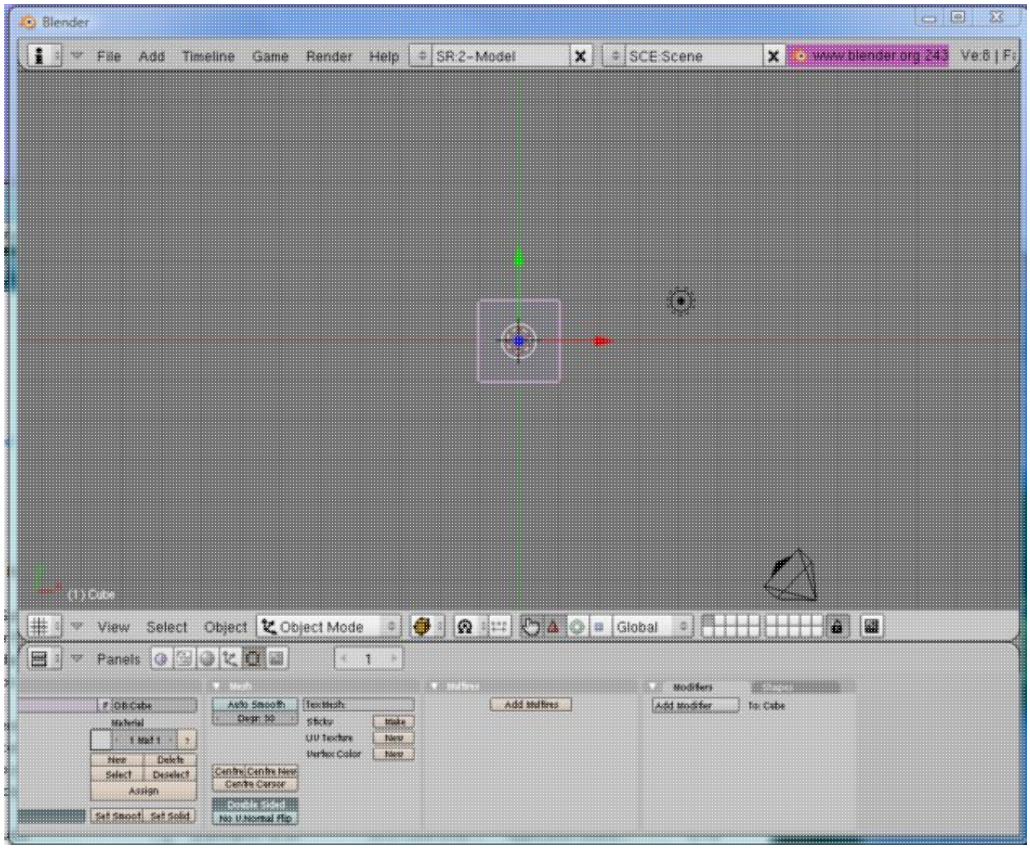


Just like the "Gus the Gingerbread Man" tutorial you see here, the BSoD Intro to Character Animation tutorial assumes no prior knowledge. It guides you through the process of making a walking, talking character from scratch and covers many powerful features of Blender not found here.

The BSoD Intro to Character Animation also has a downloadable PDF version (*3.75 MB*) for offline viewing.

Warming up

Let's start Blender. On the screen you should see, from the top view, the default set-up. A camera, a light, and a cube. The cube should already be selected, as indicated by its pink color.(*Default Blender screen as soon as you start it.*).

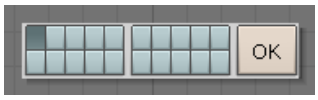


Default Blender screen as soon as you start it (small).
We will organize our working area by placing objects on different layers where we can hide them; we can also bring them back in the current scene whenever we need them. Here is how layers work:



Layer visibility controls.
Blender provides you with twenty layers to help to organize your work. You can see which layers are currently visible from the group of twenty buttons in the 3D window header (*Layer visibility controls*). You can change the visible layer with LMB and toggle visibility with â § ShiftLMB . The last layer that is turned on becomes the active layer. The active layer is where all objects that will be created are stored.

So let's clean up the place.

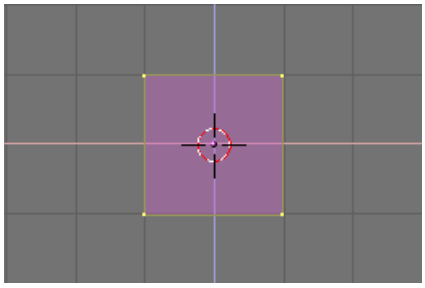


Layer control toolbox.
Select the camera and the lamp with â § ShiftRMB and press M. A small toolbox, like the one in (*Layer control toolbox*), will appear beneath your mouse, with the first button checked, which means that the selected objects are stored in Layer1. Check the rightmost button on the top row and then click on OK. This will move your camera and lamp to layer 10.

Now make sure that only Layer1 is visible, because we wouldn't want to erase the lamp or the camera ; select everything on that layer using A and erase it with X >> Erase Selected Object(s). This leaves all the room we can wish for to begin our modelling job.

Building the body

Change to the front view with 1 NumPad and add a cube --if one is not present-- by pressing Space >> Add >> Mesh >> Cube. A cube will appear. Press â Tab, and it will be in Edit Mode. (*Our cube in Edit Mode, all vertices selected*).



Our cube in Edit Mode, all vertices selected.

Edit Mode and Object Mode

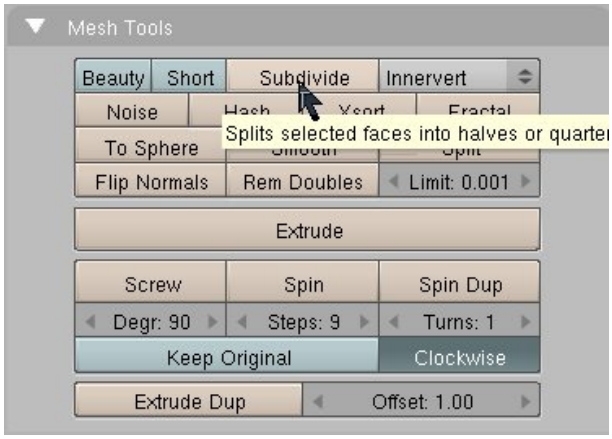
Edit Mode is a mode in which you can edit the vertices of the mesh. By default, all vertices are selected for every new object created (selected vertices are highlighted in yellow - unselected vertices are pink). In Object Mode, vertices cannot be selected or individually edited; the object can be changed only as a whole. You can press â Tab to switch between these two modes, and the current mode is indicated in the header of the 3D window.

We will call our Gingerbread man "Gus". Our first task is to build Gus's body by working on the vertices of our Cube. To see the Blender tools that we'll use for this purpose, press the button showing a square with yellow vertices in the Button window header (*The Edit Buttons Window button*), or press F9

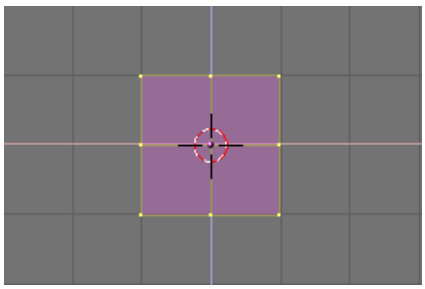


The Edit Buttons Window button.

Now locate the Subdivide button in the Mesh Tools panel and press it once (*The Mesh Tools panel in the Edit context (F9)*). This will split each side of the cube in two, creating new vertices and faces (*The cube, subdivided once*).



The Mesh Tools panel in the Edit context (F9).



The cube, subdivided once.

With your cursor hovering in the 3D window press A to deselect all elements. Vertices will turn pink.

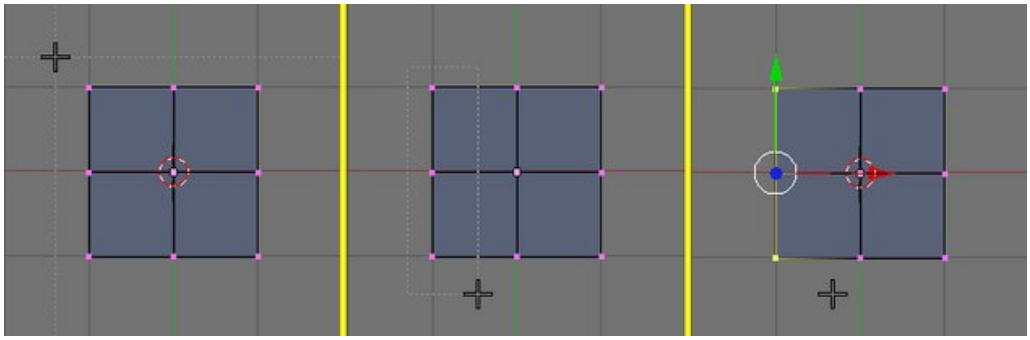


Limit selection button

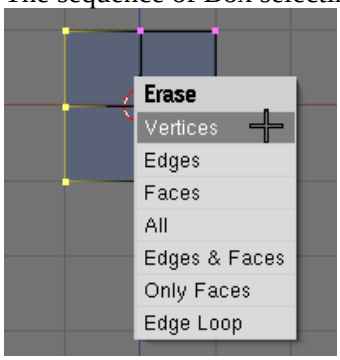
Box Select
On many occasions you may have vertices hidden behind other vertices, as is the case here. Our subdivided cube has 26 vertices, yet you can only see nine because the others are hidden. A normal **RMB** click selects only one of these stacked vertices, whereas a box select selects them all. But beware that by default this is true only for the wireframe drawtype : in any other mode, Shaded, Solid or Textured we can only select visible vertices, even with box select. To select vertices that are hidden behind others uncheck the **Limit selection button**. The button is selected in the image at the right.

You must have the **Limit selection button** *unselected*  to continue this tutorial.

Now press B, the cursor will change to a couple of orthogonal grey lines. Move the cursor above the top left corner of the cube, press and hold LMB, then drag the mouse down and to the right so that the grey box encompasses all the leftmost vertices. Now release the LMB (*The sequence of Box selecting a group of vertices*).



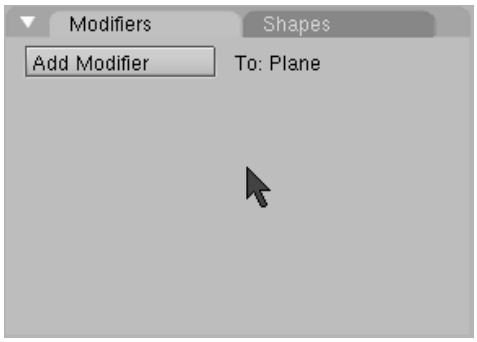
The sequence of Box selecting a group of vertices.



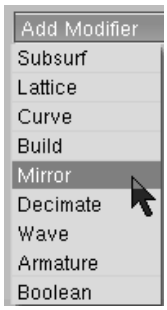
The pop-up menu of the Delete (X) action.
Press X and from the popup menu select Vertices to erase the selected vertices (*The pop-up menu of the Delete (X) action*).

Mirror modelling

To model symmetrical Objects we can use the Mirror modifier. It allows us to model only one side of Gus while Blender creates the other in real time. Go to the Edit context (F9) and find the Modifiers panel, (*The modifiers panel*).



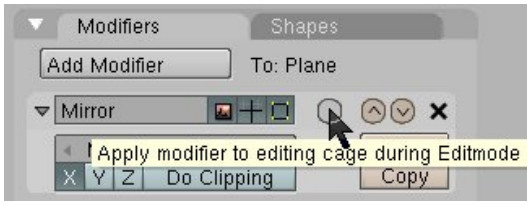
The modifiers panel.



List of modifiers.

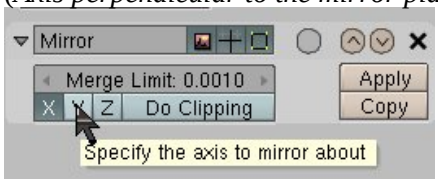
It is pretty empty for the moment. Clicking the button marked *Add Modifier* opens a list from which you'll choose Mirror(*List of modifiers*).

Nothing much seems to happen; that is because the modifiers offer quite a bit of control over what's displayed and what's not. In our case we will check the *Cage Mode* button so we can see the transparent faces in Edit Mode, (*Cage Mode button*).



Cage Mode button.

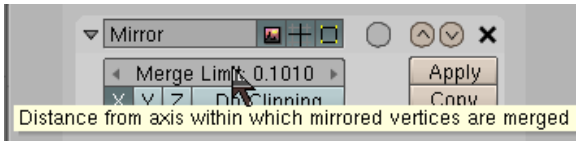
We choose the axis that will run from the modelled side of our character to the side Blender is completing by checking either the X, Y or Z button; the mirror plane is perpendicular to that axis. In our case it is the X-axis, (*Axis perpendicular to the mirror plane*).



Axis perpendicular to the mirror plane.

The Merge Limits button (*Merge Limits button*) acts as a safety net. Any vertex closer to the mirror plane, than the limit we set, will be placed exactly on the mirror plane. The limit can be set from **0.000** to **1.000** units and how big it should be depends on the nature and the scale of the current job.

For modeling Gus, a vertex that would be more than **0.1** units away from the mirror plane would be noticeable but anything closer might not. Our mesh could end up ripped in the middle if vertices that should be on the mirror plane aren't. To avoid inadvertently neglecting a wandering vertex, we should set the Merge Limits to **0.1**.



Merge Limits button.

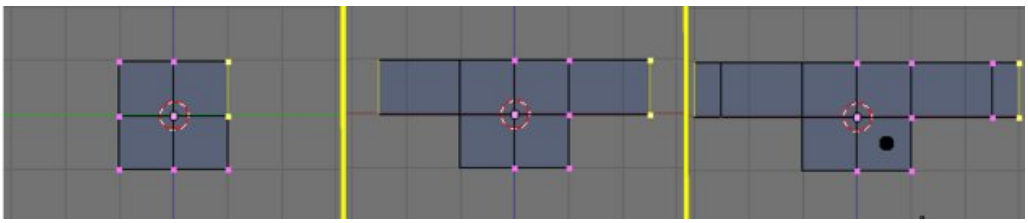
Finally, with the *Do Clipping* button checked (*Do Clipping button*), our mirror becomes a frontier that no vertex can cross. If this were to happen it would cause quite a mess. Also, when *Do Clipping* is active, every vertex that is **on** the mirror sticks to it.



Do Clipping button.
As you can see, the Mirror modifier gives us a lot of features to make our lives easier.

Arms and Legs

Let's create Gus's arms and legs. Using the sequence you just learned, Box Select the two top-right-most vertices (*Extruding the arm in two steps*), left) which will actually select the other two behind them, for a total of four vertices. Press E and click on the Region menu entry to extrude them. This will create new movable vertices and faces which you can move with the mouse. Move them one and a half squares to the right, then click LMB to fix their position. Extrude again with E then move the new vertices another half a square to the right. (*Extruding the arm in two steps*) shows this sequence.



Extruding the arm in two steps.

Undo/Redo

Blender has two Undo features, one for Edit Mode and the other for Object Mode.

In Edit Mode press CtrlZ to Undo and keep pressing CtrlZ to roll back changes as long as the Undo buffer will allow; â § ShiftCtrlZ re-does changes. AltU opens a menu with a list of possible undos so that you can easily find the point you want to revert to.

Two things to remember:

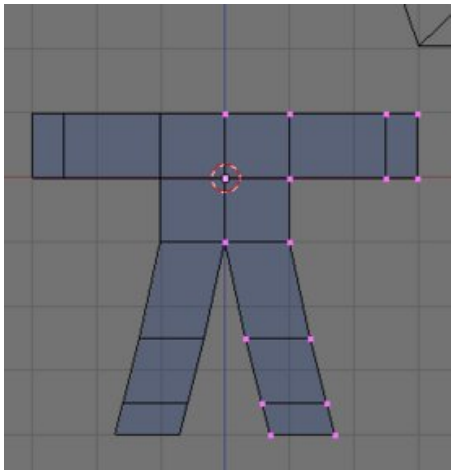
- Undo in Edit Mode works only for the Object currently in that mode.
- Undo data is not lost when you switch out of Edit Mode, but it is as soon as you start editing a *different* Object in Edit Mode.

In Object Mode the same shortcuts apply. CtrlZ to undo, â § ShiftCtrlZ to redo and AltU to see the history. If you made changes in Edit Mode that are not lost for that Object, they will all be undone in one single shot with CtrlZ when this step --marked as Edit Mode in the Object Mode (AltU) history-- has its turn.

If you change your mind in the middle of an action, you can cancel it immediately and revert to the previous state by pressing Esc or RMB.

Coincident vertices

Extruding works by first creating new vertices and then moving them. If in the process of moving you change your mind and press Esc or RMB to cancel, the new vertices will still be there, on top of the original ones! The simplest way to go back to the state before you started extruding is to Undo (CtrlZ). It is sometimes useful to intentionally create new vertices this way and then move, scale or rotate them by pressing G,S or R.



Body.

Gus should now have a left arm that you modelled (he's facing us) and a right arm that Blender added. We will build the left leg the same way by extruding the lower vertices three times. Try to produce something like in (*Body*). If you are using *Extrude - Region*, you will have to clear the transformation constraint *Mirroring* by unchecking *Do Clipping* in the Mirror Modifier panel, and rechecking it after extruding (otherwise Gus will end up with a skirt rather than pants).

You can free the movement of the extruding vertices by clicking MMB after you have pushed E but before you click LMB. If you do not do this your legs will end up going straight down, rather than down and to the side as in (*Body*).

We're done with mirror modelling. In the next steps we will experiment with other techniques. We need to make the right part of our model *real* since nothing done with modifiers is permanent unless we *apply* the changes. With Gus being in Object Mode (press â Tab), click on the *Apply* button of the Mirror modifier.

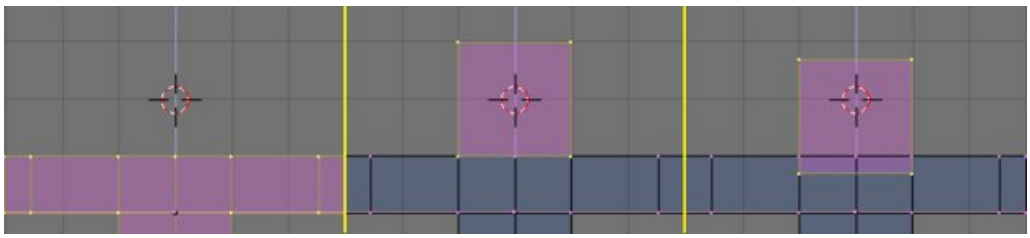
The Head

Gus needs a head.

Change back to Edit Mode (press â Tab)

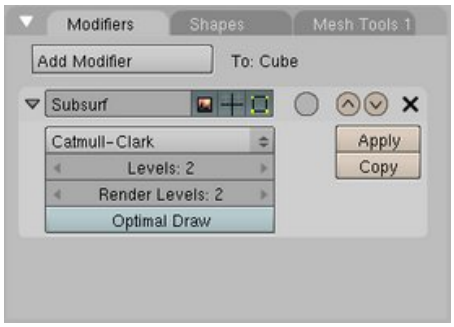
Move the cursor to exactly one square above Gus's body (leftmost image of *Adding the head*) and add a new cube (Space>ADD>Cube). To place the cursor at a specific grid point, position it next to where you want it and press **SHIFT+S** to bring up the Snap Menu. **Cursor to Grid** places the cursor exactly on a grid point. That's what we want right now. **Cursor to Selection** places it exactly on the selected object, which is sometimes handy.

Now press G to switch to Grab Mode and move the newly created vertices down, constraining the movement by moving the head down a bit and clicking MMB, for about one third of a grid unit (rightmost image of *Adding the head.*).



Adding the head.

SubSurfaces *(subsurf)*



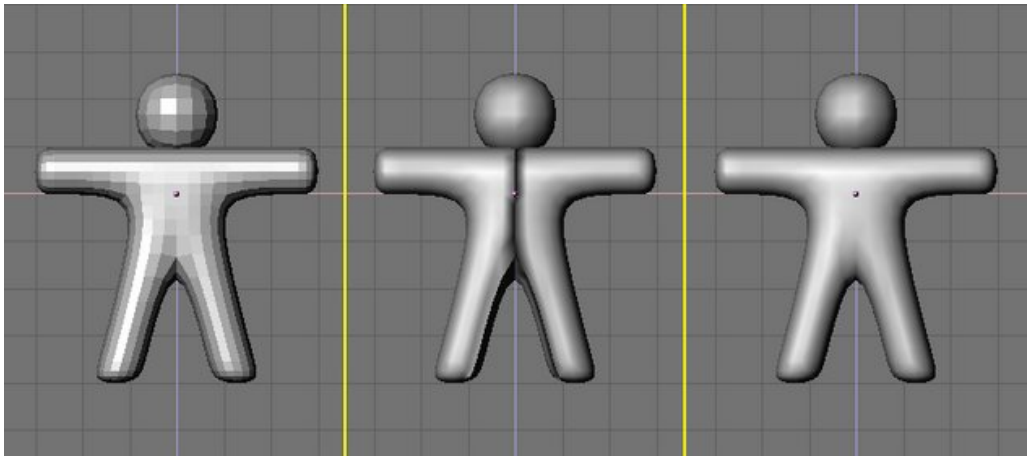
The Subsurf modifier in the Modifiers panel of the Editing context (F9)

For the next step we'll need to select all of Gus, and not just his head. So far what we have produced is a rough figure at best. To make it smoother, locate the Modifier panel in the *Editing* context (F9) and add a Subsurf modifier, (*The* Subsurf*modifier in the Modifiers panel*). Be sure to set both *Levels* NumButtons below or at 2. The first Level is for what you'll see in the 3D Window area, the second for the renderer.

SubSurfaces

SubSurfacing is an advanced modelling tool, it dynamically refines a given coarse mesh creating a much denser mesh and locating the vertices of the finer mesh so that they smoothly follow the original coarse mesh. The shape of the Object is still controlled by the location of the coarse mesh vertices, but the rendered shape is a finely smooth mesh.

- Switch out of Edit Mode (â Tab) and from the current Wireframe mode to Solid mode using Z to have a look at Gus. He should look like (*Setting Gus to smooth*, left).



Setting Gus to smooth.

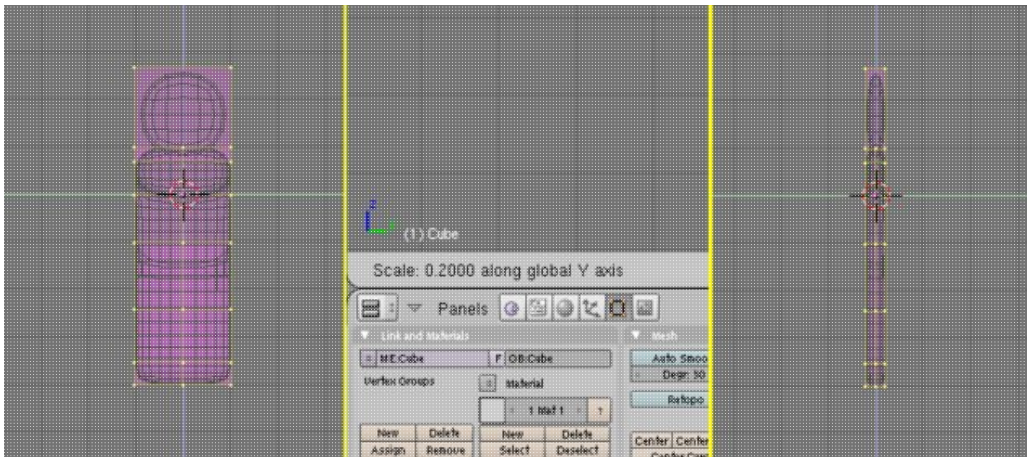
- To make Gus look smooth, press the SetSmooth button found in the Link and Material panel of the Editing context (F9). Gus will now appear smooth although he may wear some funny black lines in his middle. This is usually avoided if you used the Mirror Modifier but it might happen when extruding and flipping, as it was done before the modifier was introduced. (*Setting Gus to smooth.*, middle). These lines appear because the SubSurf's finer mesh is computed using information about the coarse mesh normal directions, which may not all point in the right direction, that is, some face normals might point outward and some inward. To reset the normals, switch back to Edit Mode (â

Tab), select all vertices (A), and press CtrlN. Click with LMB on the Recalculate normals outside box which appears. Now Gus should be nice and smooth (*Setting Gus to smooth*, right).

Press MMB and drag the mouse around to view Gus from all angles. Oops, he is too thick!

Constrained Scaling

(*Slimming Gus using constrained scaling.*, left).



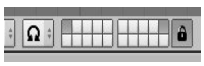
Slimming Gus using constrained scaling.
Let's make Gus thinner:

- Switch to Edit Mode if you are not there already (â Tab), then back to Wireframe mode. (Z), Switch to side view using NUM3 and select all vertices with A. You can do the following steps just as well in Object Mode, if you like.
- Press S and start to move the mouse horizontally. (Click MMB to constrain scaling to just one axis or press Y to obtain the same result). If you now move the mouse toward Gus he should become thinner but remain the same height.
- The header of the 3DWindow toolbar shows the scaling factor (*Slimming Gus using constrained scaling.*, center). Press and hold Ctrl. The scale factor will now vary in discrete steps of value 0.1. Scale Gus down so that the factor is 0.2, then set this dimension by clicking LMB, (*Slimming Gus using constrained scaling.*, right).
- Return to Front view and to Solid mode (Z), then rotate your view via MMB. Gus is much better now!

Let's see what Gus looks like

HI just about ready to see our first rendering, but first, we have some work to do.

- Switch to Object Mode if not already there (â Tab).

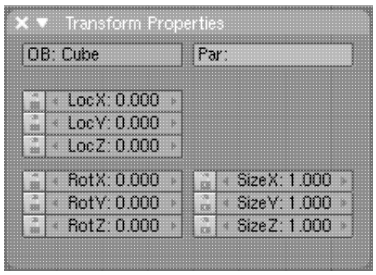


Making both layer 1 and 10 visible.

- â § ShiftLMB on the top right small button of the layer visibility buttons in the 3DWindow toolbar (*Making both layer 1 and 10 visible.*) to make both Layer 1 (Gus's layer) and Layer 10 (the layer with the camera and the lamp) visible.

A Tip
Remember that the last layer selected is the active layer, so all subsequent additions will automatically be on layer 10.

- Press N to bring up the Transform Properties window (*The Panel for numerical input of object position/rotation etc*). The location of the camera is specified by LocX, LocY, and LocZ.



The Panel for numerical input of object position/rotation etc.

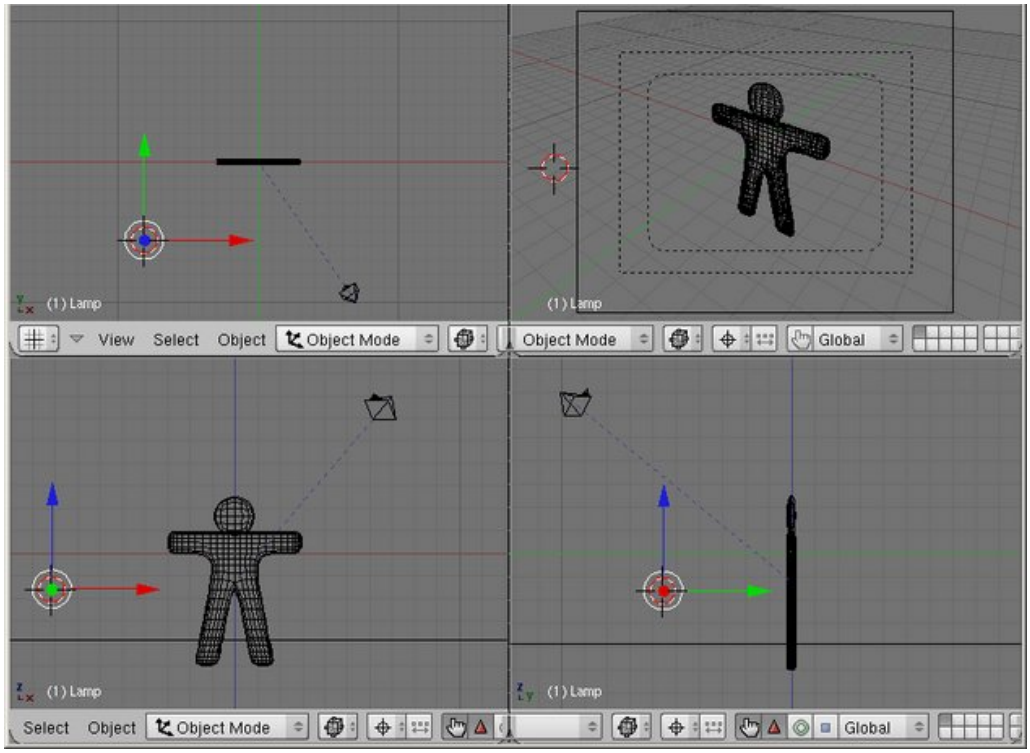
- Select the camera (RMB) and move it to a location like (x=7, y=-10, z=7). Do this by pressing G and dragging the camera. You may need to change views and move the camera a second time to adjust all three coordinates. If you prefer to enter numerical values for an object's location you can do so by holding â § Shift and clicking LMB on a NumButton and then entering the desired value. Remember to press â µ Enter to confirm your input.

Camera setup

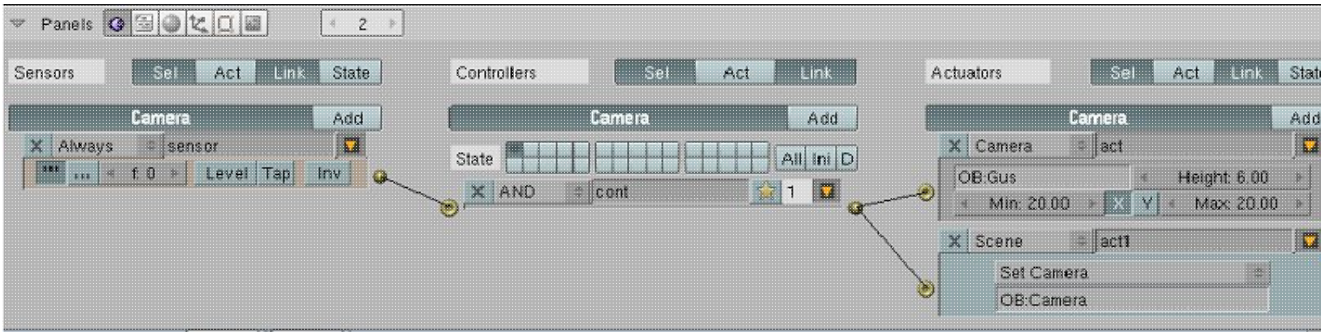
To make the camera point at Gus, keep your camera selected then select Gus via â § ShiftRMB. The camera should be magenta and Gus light pink. Now press CtrlT and select the TrackTo Constraint entry in the pop up. This will force the camera to track Gus and always point at him. This means that you can move the camera wherever you want and be sure that Gus will always be in the center of the camera's view.

Tracking
If you choose the option **Old Track** and the camera has a rotation of its own, as is often the case, it could point in an unexpected direction. In that case select the tracking object (in our example the camera), and press **ALT-R** to remove the object's rotation. Once you do this the camera will really track Gus.

(*Camera position with respect to Gus*) shows top, front, side and camera view of Gus. To obtain a camera view press 0 NumPad or select View>>Camera.



Camera position with respect to Gus.
if tracking doesnt work in game mode then you can track gus with logic go to logic button



make sure camera OB:Gus the object you want to link the camera to and scene, set camera is selected with the OB:Camera the camera that want to use.

The Ground

Now we need to create the ground for Gus to stand on.

- In top view (7 NumPad or View>>Top), and in Object Mode, add a plane

(Space>>Add>>Mesh>>Plane).

Note
It is important to be out of Edit Mode, otherwise the newly added object would be part of the object currently in Edit Mode, as when we added Gus' head.

- Switch to the Front view (1 NumPad or View>>Front) and move (G) the plane down to Gus's feet, using Ctrl to keep it aligned with Gus.

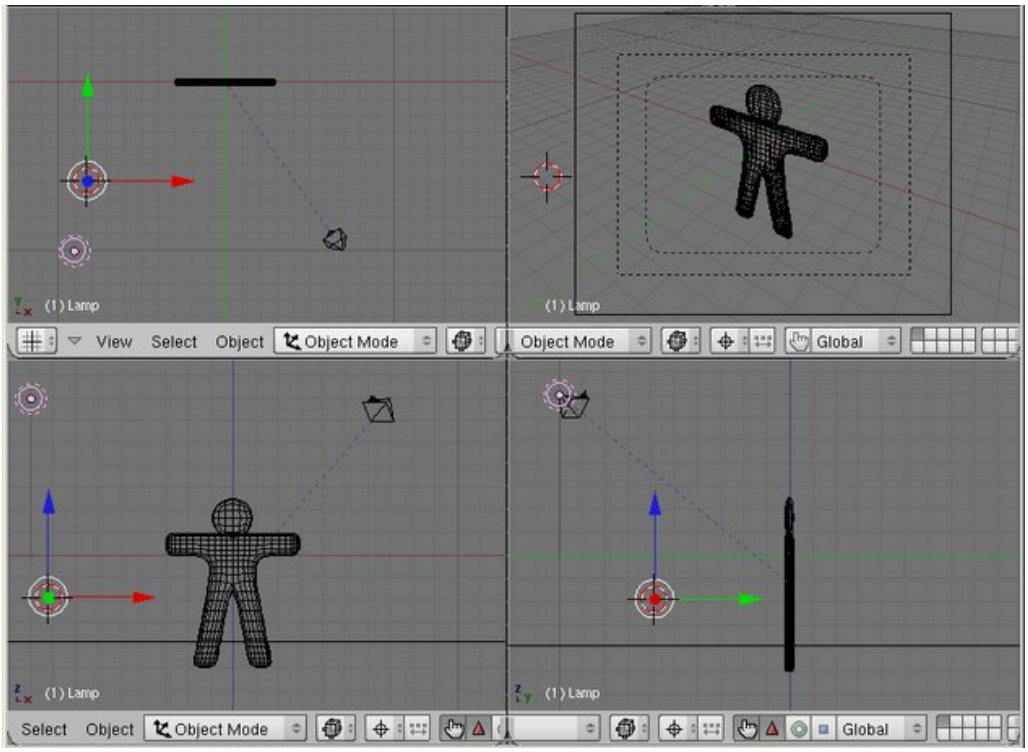
- Go to Camera view (0 NumPad or View>>Camera) and, with the plane still selected, press S to start scaling.

- Enlarge the plane so that its edges extend beyond the camera viewing area, as indicated by the outer black dashed rectangle in Camera view.

Lights

Now, lets add some light!

- In Top view (7 NumPad), move the existing Lamp light (if you do not have a Lamp light in your scene you can add one with Space>>Add>>Lamp>>Lamp) in front of Gus, but on the other side of the camera; for example to (x= -9, y= -10, z=7) (*Inserting a Lamp.*).



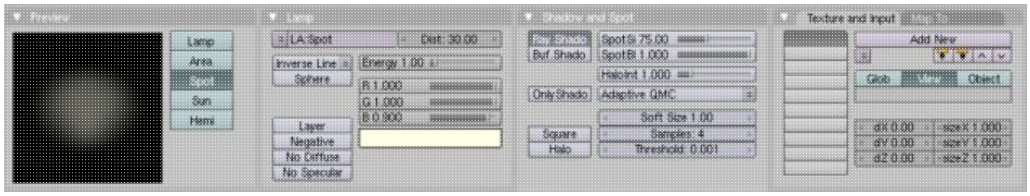
Inserting a Lamp.



The Lamp buttons window button.

- Switch to the Shading context (F5) and then the Lamp buttons window via the sub-context button with a lamp in the Button Window header (*The Lamp buttons window button.*).

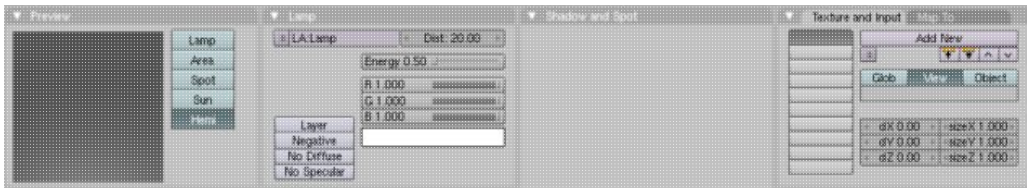
- In the Buttons Window, Preview Panel, press the Spot toggle button to make the lamp a Spotlight (*Spot light settings.*) of pale yellow (R=1, G=1, B=0.9). Adjust Samples: to 4 and SpotBl: to 1.0.



Spot light settings.

- Make this spotlight track Gus just as you did for the camera by selecting Spot, â § Shift, then Gus, then by pressing CtrlT>>TrackTo Constraint.

- Add a second lamp that provides more uniform fill light via (Space>>Add>>Lamp>>Hemi). Set its Energy to 0.5 (*Hemi lamp settings*). Move it a little above the camera (x= 7, y= -10, z=9) and set it to track Gus as before.



Hemi lamp settings

Two lamps?

Use two or more lamps to help produce soft, realistic lighting, because in reality natural light never comes from a single point.

Rendering

We're almost ready to render. As a first step, press the Scene context button and the Render sub-context button in the Button window header (*The Rendering buttons window buttons.*).



The Rendering buttons window buttons.

We will use the default rendering settings, as shown in (*The Rendering Buttons window*).



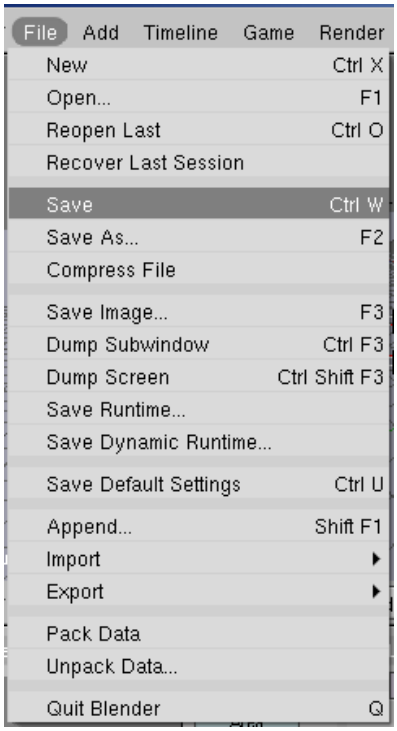
The Rendering Buttons window

Now press the RENDER button or F12. The result, shown in (*Your first rendering. Congratulations!*), is actually quite poor. We still need materials, and lots of details, such as eyes, and so on.



Your first rendering. Congratulations!

Saving our work



The Save menu.

If you have not done so already, now would be a good time to save your work, via the File>>Save menu shown in *The Save menu.*, or CtrlW. Blender will warn you if you try to overwrite an existing file.

Blender does automatic saves into your system's temporary directory. By default, this happens every four

minutes and the file name is a number. Loading these saves is another way to undo unwanted changes.

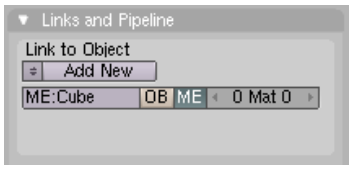
Materials and Textures

It's time to give Gus some nice cookie-like material.



The Material Buttons window Button.

- Select Gus. Then, in the Button Window header, select the Shading Context by pressing the red dot button (*The Material Buttons window Button.*) or pressing F5. Then press the red dot sub-context button to access the Material panels.

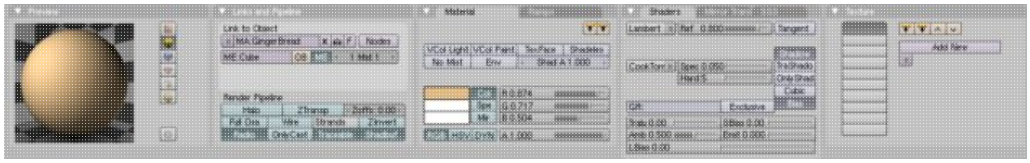


The Material Menu button.

- The Button window will be almost empty because Gus has no materials yet. To add a material, click on the Menu Button in the Material Panel (the one with two triangles, pointing up and down) and select Add New (*The Material Menu button.*).

- The Buttons window will be populated by Panels and Buttons and a string holding the Material name, something like "Material.001", will appear next to the white square button. Click the name and change it to something meaningful, like "GingerBread" (don't type the quotes).

- Modify the default values as per (*The Material Buttons window and a first gingerbread material*) to obtain a first rough material. Note that you must click the Shaders tab to reveal the shader panel.



The Material Buttons window and a first gingerbread material.

- Press the Menu Button in the Textures Panel area (*The Textures menu button in the Material Buttons*) and select Add new. We're adding a texture in the first channel. Call it "GingerTex."



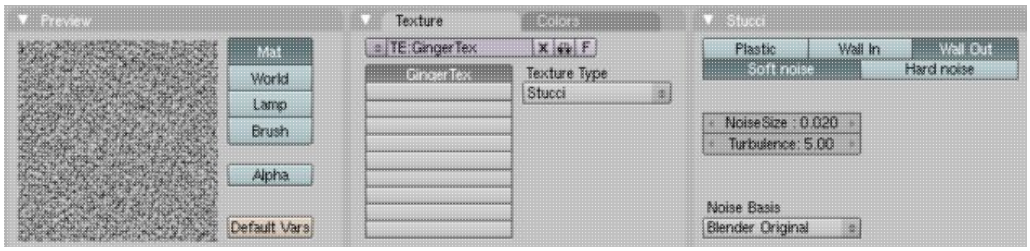
The Textures menu button in the Material Buttons



The Texture Buttons window Button.

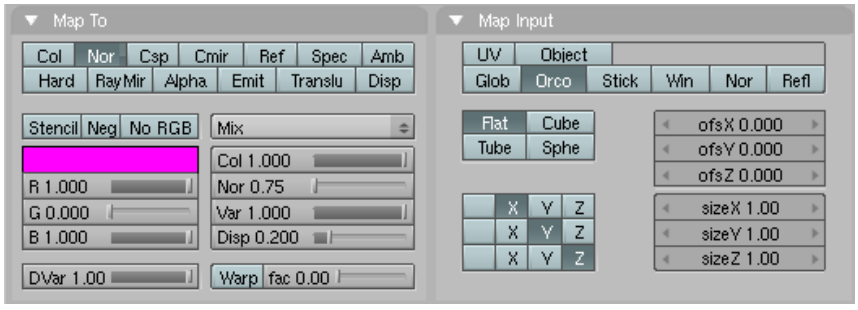
- Select the Texture Buttons by clicking the button in (*The Texture Buttons window Button*) or by pressing F6.

- From the columns of ToggleButtons which appear in the Texture panel select Stucci and set all parameters as in (*The Texture Buttons window with a stucci texture*).



The Texture Buttons window with a stucci texture.

- Return to the Material buttons (F5) and set the Map Input and Map To tabs of the Texture Panel as in (*Settings for the Stucci texture in the Material Buttons window*). Release the Col Toggle Button and set the Nor Toggle Button, then raise the Nor slider to **0.75**. These changes will make our Stucci texture act as a "bumpmap" and make Gus look more biscuit-like.



Settings for the Stucci texture in the Material Buttons window.

- Now add a second texture, name it "Grain", and make it affect only the Ref property with a **0.4** Var (*Settings for an additional Noise texture in channel 2*). The texture itself is a plain Noise texture.



Settings for an additional Noise texture in channel 2.

- Give the ground an appropriate material, such as the dark blue one shown in (*A very simple material for the ground*).



A very simple material for the ground.

Eyes and detail

To give some finishing touches we'll add eyes and some other details.

- First make Layer 1 the only one visible by clicking with LMB on the layer 1 button (*Layer visibility buttons on toolbar*). This will hide the lamps, camera, and ground.

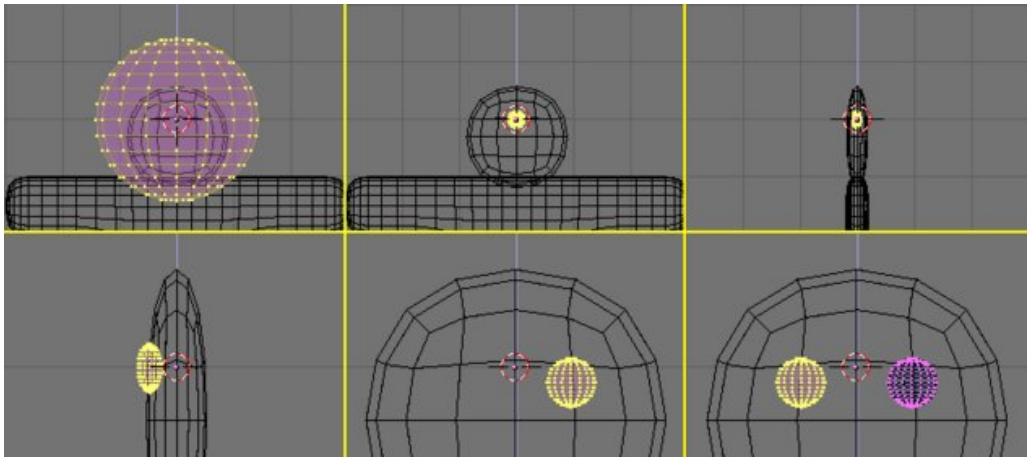


Layer visibility buttons on toolbar.

- Place the cursor at the center of Gus's head. (Remember that you are in 3D so be sure to check at least two views to be sure!)

- In Object Mode, add a sphere (Space>>ADD>>Mesh>>UVsphere). You will be asked for the number of Segments: (meridians) and Rings: (parallels) into which to divide the sphere. The default of **32** is more than we need here, so use a value of **16** for both. The sphere is in the first image at the top left of the sequence in (*Sequence for creation of the eyes*).

- Scale the sphere down (S) to a factor of about **0.15** in all dimensions, then switch to side view (3 NumPad) and scale it only in the horizontal direction (Y) a further **0.5**, see the second two images in (*Sequence for creation of the eyes*).



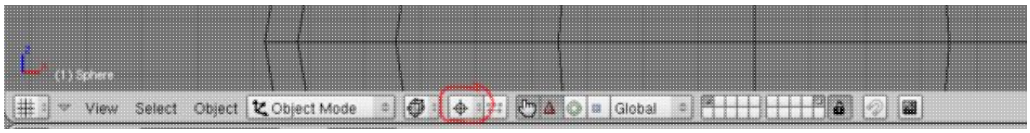
Sequence for creation of the eyes.

- Zoom a little if necessary via + NumPad, Wheel, or CtrlMMB, and drag the sphere (G) to the left so that it is halfway into the head, as shown in the first image in the second row of (*Sequence for creation of the eyes*).

- Return to front view (1 NumPad) and move the sphere sideways, to the right. Place it where Gus should have an eye.

Flipping a duplicate around the cursor

Switch to Edit Mode (â Tab). Select the crosshair pivot button in the header of the 3D window (*pivot: 3D Cursor*). With just the eye selected, change to Edit Mode. The eye should still be selected (if not, press A to select all), now press â § ShiftD to duplicate and Esc to stop placing it with the mouse. Then press M to mirror, X to mirror around the X axis, followed by LMB or â µ Enter to confirm the mirror. Return the pivot button to its default setting (Median Point).



The crosshair pivot button.

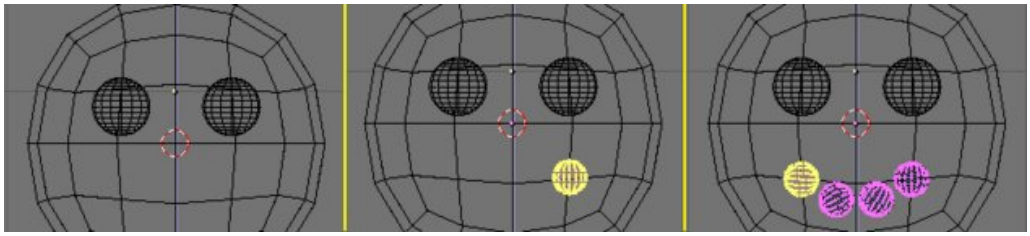
Mirroring

Mirroring can also be done in object mode using Ctrl M.

Now Gus has two eyes.

Mouth

- Exit Edit Mode (â Tab), and place the cursor as close as you can (remember the â § ShiftS key) to the center of Gus's face. Add a new sphere and scale and move it exactly as done before for the eyes, except make its over scale smaller (0.1 instead of 0.15). Place it below and to the right of the cursor, centered on the SubSurfed mesh vertex, as shown in the middle frame in the image below: *Creating a mouth with Spinning tools.*.



Creating a mouth with Spinning tools.



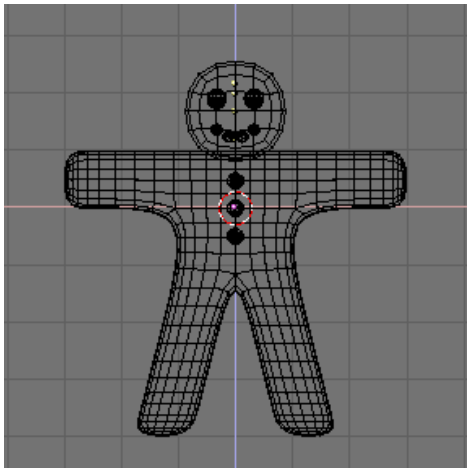
The Spin Tools buttons in the Edit Buttons window.

- Switch to Edit Mode (â Tab). Now, in the Edit Buttons (F9), locate the group of buttons at bottom in the Mesh Tools Panel (*The Spin Tools buttons in the Edit Buttons window*.). Set Degr: to **90**, Steps: to **3**, and verify that the Clockwise: TogButton is on. Then, with all vertices still selected, press SpinDup. This will create three duplicates of the selected vertices on an arc of **90** degrees, centered around the cursor. The result should be Gus's mouth, like the last image of the sequence shown in *Creating a mouth with Spinning tools.*.

- Now go back to Object Mode and add three more spheres (below the head and centered along the Z-axis) to form Gus's buttons. Once you have made one button, you can simply exit Edit Mode, press â § ShiftD to create a duplicate, and move the duplicate into place, as shown in *The complete Gus!*.

Attaching the spheres

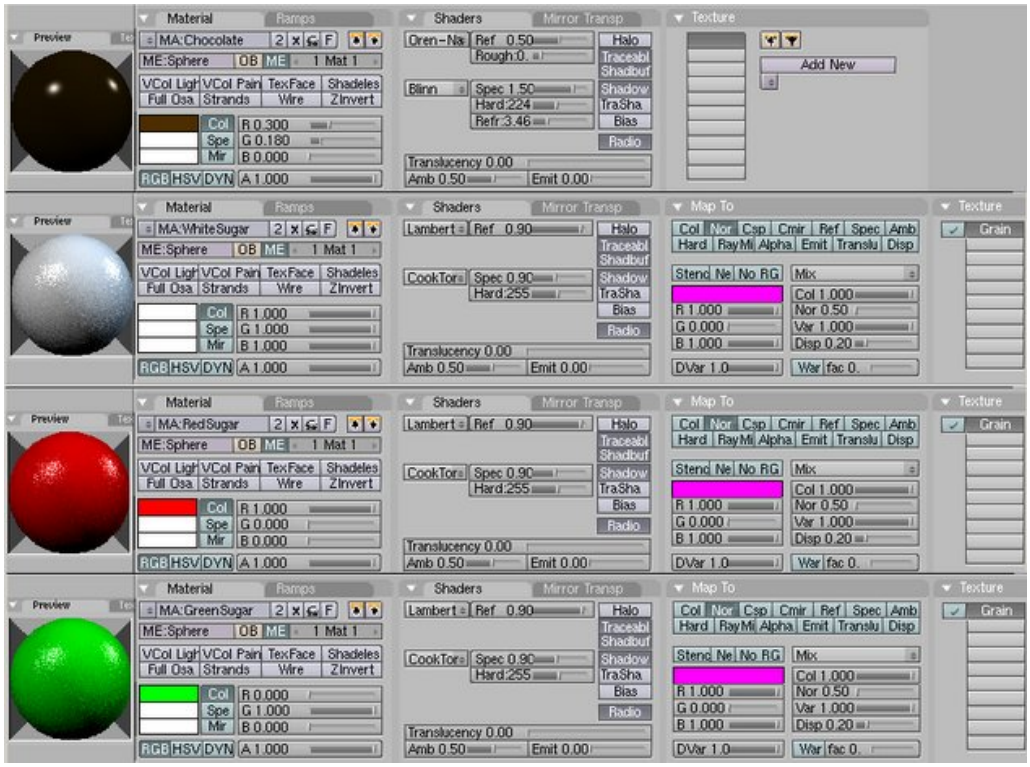
If we want to be able to grab Gus and move him around as a whole (this goes beyond the animation in the second part of this tutorial), we now need to attach the small spheres representing eyes, mouth, and buttons to the body. Enter Object Mode and press A until nothing is selected. Now right click one sphere (if more than one is selected as a group, that's ok). Holding â § Shift, select the body. Then hit CtrlP and left click Make parent on the pop up. Deselect everything and repeat the process to attach each element.



The complete Gus!

Eyes material

Give the eyes a chocolate-like material, like the one shown at the top in *Some other candy materials*.. Give the mouth a white sugar-like material, like the second one shown in (*Some other candy materials*), and give the buttons a red, white, and green sugar-like material. These are shown from top to bottom in (*Some other candy materials*) too.



Some other candy materials.



Material Menu.
Objects sharing a material
To give one object the same material as another object, select that material in the **Material Menu** list which appears when you press the Menu Button ButtonWindow **Material** Panel, see (*Material Menu*).

Rendering

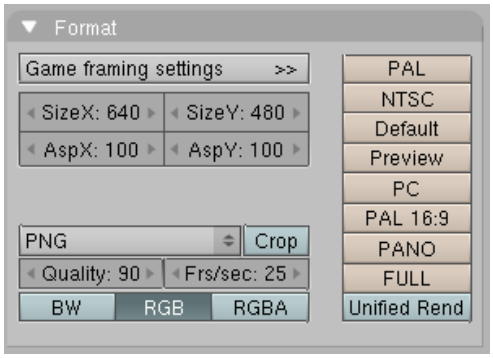
Once you have finished assigning materials, make layer **10** visible again (remember how? Hint, look at the 3D window header), so that lights and the camera also appear, and do a new rendering (F12).

The result should look more or less like (*The complete Gus still rendering*).



The complete Gus still rendering.

Saving



File type selection menu in the Rendering Buttons window.
Save your image by pressing F3. Enter the name of your image in the file window and save.

You must choose the image format (JPEG, PNG, and so on) by setting it in the Rendering buttons *before* pressing F3 (*The Rendering buttons window buttons*) and using the Menu (*File type selection menu in the Rendering Buttons window*) in the Format Panel.

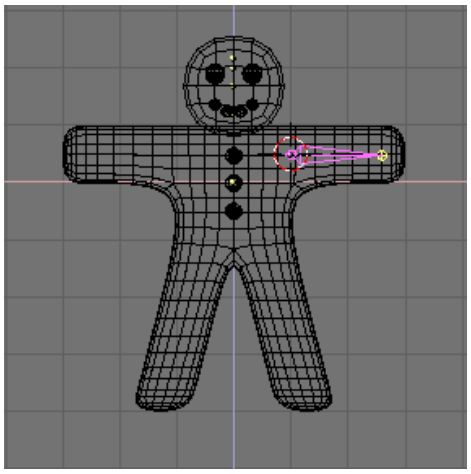
Blender does *not* add an extension to the file name; you must enter one if you wish.

Your First Animation in 30 plus 30 Minutes Part II

If we were going for a still picture, our work up to this point would be enough, but we want Gus to move! The next step is to give him a skeleton, or Armature, which will move him. This is called the fine art of rigging. Gus will have a very simple rigging: four limbs (two arms and two legs) and a few joints (no elbows, only knees), but no feet or hands.

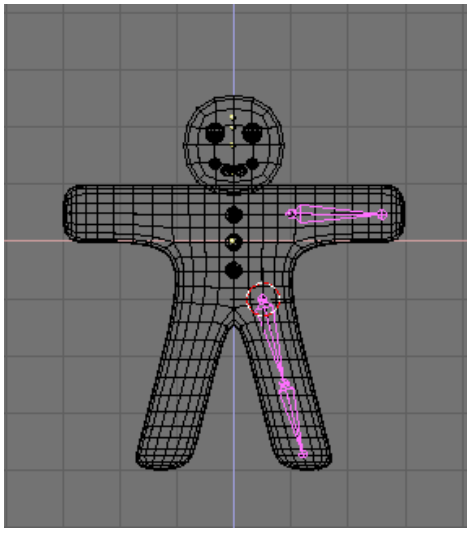
Rigging

To add the rigging:



Adding the first bone, an elbowless arm.

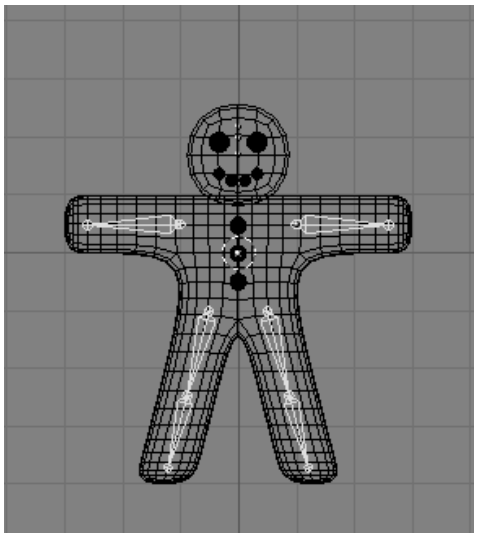
- In Object mode, set your 3D cursor where Gus's shoulder is, and press Space>>Add>>Armature. A rhomboidal object will appear, which is a bone of the armature system. Enter Edit mode. The end of the bone is selected (yellow).
- Now in Edit mode, place the other end of the armature in Gus's hand by grabbing (G) and moving the end, (*Adding the first bone, an elbowless arm*). We don't need any other bones right now. You should now have one bone running from the shoulder to the hand area. As you move the end, you will notice that the whole bone gets bigger; you really are scaling up the bone.



Adding the second and third bones, a leg bone chain.

- Stay in Edit mode, then move the cursor to where the hip joint will be and add a new bone (Space>>Add>>Bone).
- Grab (G) and move the yellow end of the new bone to the knee area.
- Now "chain" a new bone from the knee to the foot by CtrlLMB in the area of the foot. A new "chained" bone will appear automatically linked with the knee and ending at the foot, (*Adding the second and third bones, a leg bone chain*). Another way of "chaining" the new bone would be to extrude using the (E). This variation creates the new bone and places you in grab mode automatically. Use the LMB to then place the bone.

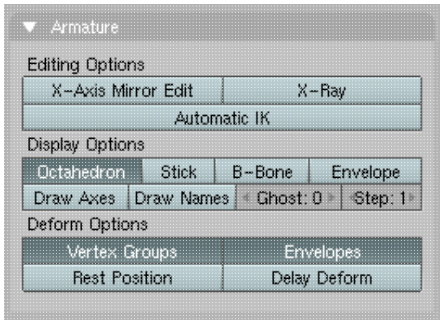
Bone position
The bones we are adding will deform Gus's body mesh. To produce a neat result, try to place the bone joints as shown in the illustrations. We now have three bones that make up Gus's Armature.



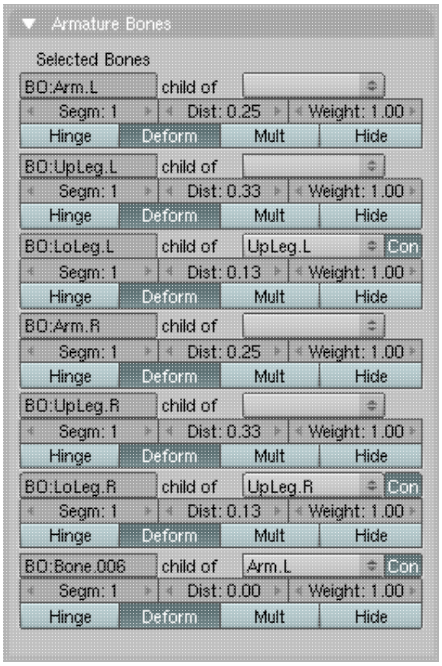
The complete armature after duplicating and flipping.

- Now place the cursor in the center and select all bones with A. Duplicate them with â § ShiftD and exit grab mode with Esc. Make sure the cursor is selected as the rotation/scaling pivot. Flip the bones along the X axis relative to the cursor with CtrlM and then X. You end up with (*The complete armature after duplicating and flipping*).

Once you've selected all of the bones (A), the Edit Buttons window should show an Armature Panel and Armature Bones Panel which contains the Armature controls (*Armature panel* and *Armature Bones panel*).



Armature panel.



Armature Bones panel.

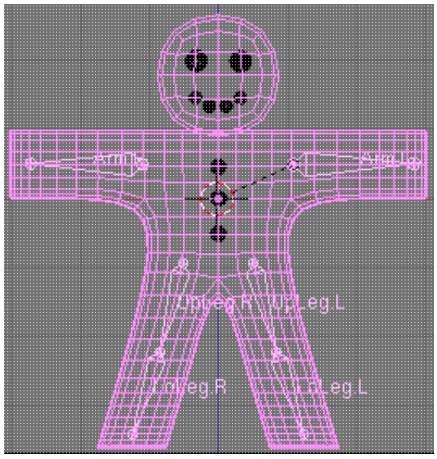
Press the Draw Names button to see the names of the bones in 3D View, then LMB click on the names in the Edit Button window to change them to something appropriate like Arm.R, Arm.L, UpLeg.R, LoLeg.R, UpLeg.L and LoLeg.L, see (*The Edit Buttons window for an armature*). Exit EditMode with (â Tab).

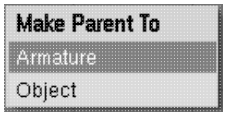
Naming Bones
It is very important to name your bones with a trailing '.L' or '.R' to distinguish between left and right ones, so that the Action editor will be able to automatically flip your poses.

Skinning

Now we must make it such that a deformation in the armature causes a matching deformation in the body. We do this with *Skinning*, which assigns vertices to bones so that the former are subject to the latter's movements.

- In ObjectMode, select Gus's body, then â § Shift select the armature so that the body is magenta and the armature is light pink.
- Now we need to parent the body to the armature. That is achieved by doing Space>>Object>>Parent>>Make Parent (or, alternatively, press CtrlP). The (*Parenting menu*) will appear. Select the Armature entry.





Parenting menu.



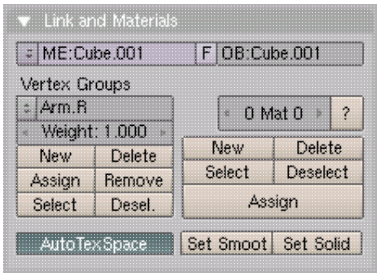
Create Vertex Group menu.

Armature parented.

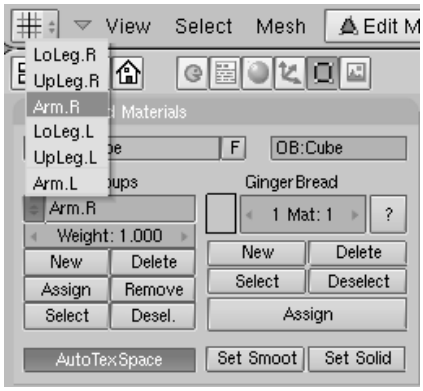
- A new menu appears, asking if you want Blender to do nothing, create empty vertex groups, or create and populate vertex groups (*Create Vertex Group menu*). The last option is considered *automatic skinning*.

- We'll use the automatic skinning option. Go ahead and select Create From Closest Bones. Now select just Gus's body and switch to EditMode (â Tab). Notice in the Edit Buttons Window (F9) the presence of the "Vertex Groups" menu and buttons in the Link and Materials Panel, (*The vertex groups buttons in the Edit Buttons window*).

Update Note
With Blender v2.46, the option **Create From Closest Bones** is no longer available, use **Create From Bone Heat** instead.

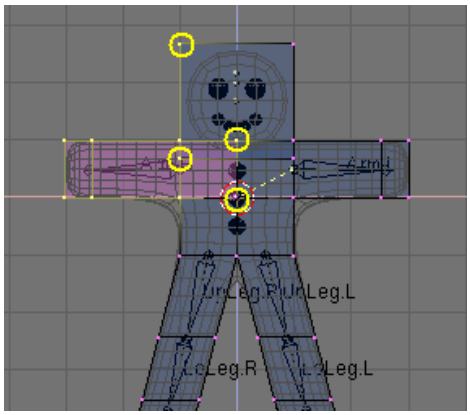


The vertex groups buttons in the Edit Buttons window



The menu with the vertex groups automatically created in the skinning process.

By pressing the Menu Button a menu, with all available vertex groups, pops up --six in our case. But a truly complex character, with hands and feet completely rigged, can have tens of them! See (*The menu with the vertex groups automatically created in the skinning process*). The buttons Select and Deselect show you which vertices belong to which group.



Gus in EditMode with all the vertices of group Arm.R selected.
Select the Right arm group (**Arm.R**) and, with all vertices de-selected (A, if needed) press Select. You should see something like (*Gus in EditMode with all the vertices of group Arm.R selected*).

If you don't see the same thing then you probably placed the bones in just the right place such that the *auto skinning* process did a better job of matching vertices with bones. It is highly unlikely that the skinning

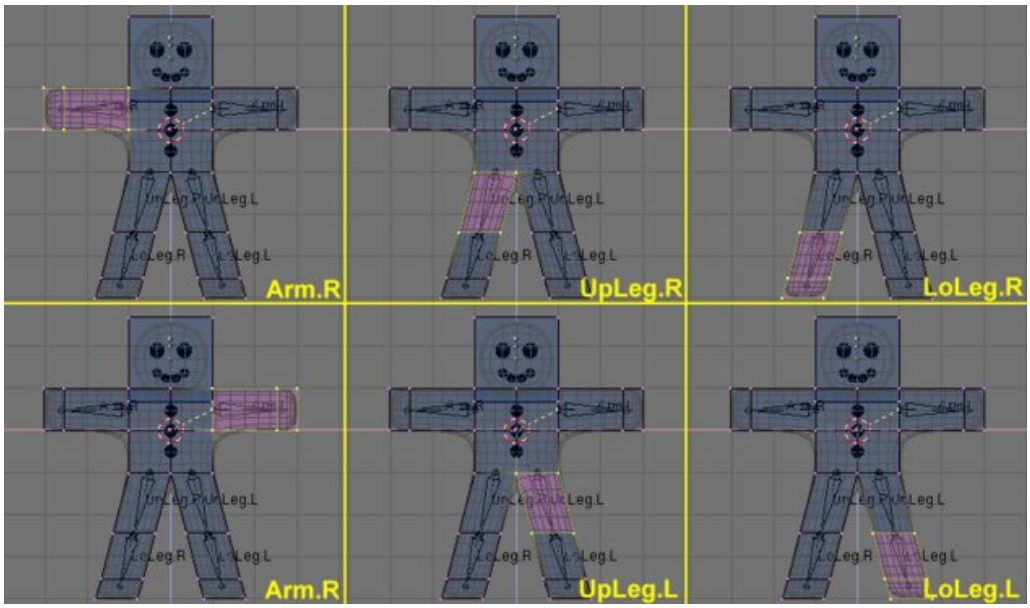
process matched the vertices to the bones as exactly as you may expect. This requires that you begin to manually adjust the grouping as described in the following sections.

The vertices marked with yellow circles in (*Gus in EditMode with all the vertices of group Arm.R selected*) belong to the deformation group, however, they should not.

The *auto skinning* process found that they were very close to the bone so it added them to the deformation group. We don't want them in this group since some are in the head and some are in the chest, adding them to the deformation group would deform those body parts as well.

To remove them from the group, deselect all the other vertices, those which should *remain* in the group using Box selection (B), but use MMB, not LMB, to define the box, so that all vertices within the box become deselected.

Once the 'undesired' vertices are selected, press the Remove button (*The vertex groups buttons in the Edit Buttons window*) to eliminate them from group (**Arm.R**). Deselect all (A) then check another group. Check them all and be sure that they look like those in (*The six vertex groups*).



The six vertex groups.

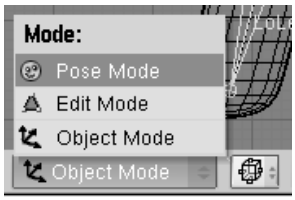
Vertex groups

Be very careful when assigning or removing vertices from vertex groups. If later on you see unexpected deformations, ME might have forgotten some vertices, or placed too many in the group. You can modify your vertex groups at any time.

Other details
Our deformations will affect only Gus's body, not his eyes, mouth, or buttons, which are separate objects. While this is not an issue to consider in this simple animation, it's one that must be taken into account for more complex projects, for example by parenting or otherwise joining the various parts to the body to make a single mesh. (We'll describe all of these options in detail in later Chapters).

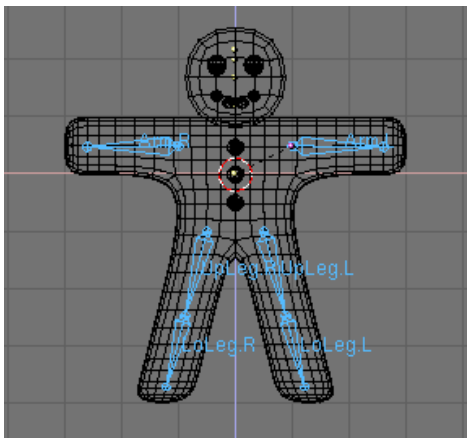
Posing

Once you have a rigged and skinned Gus you can start playing with him as if he were a doll, moving his bones and viewing the results.



Mode menu in the 3D Window header.

- Select the armature only, then select Pose Mode from the "Mode" Menu (*Mode menu in the 3D Window header*). This option only appears if an armature is selected.
- The armature will turn blue. You are now in Pose Mode. If you now select a bone it will turn cyan, not pink, and if you move it (G), or rotate it (R), the body will deform!



You are in pose mode now!

Original position

Blender remembers the original position of the bones. You can set your armature back by pressing AltR to clear the rotation and AltG to clear the location. Alternatively, the Rest Position button may be used to temporarily show the original position.

Inverse Kinematics

Inverse Kinematics (IK) is where you actually define the position of the *last* bone in the chain, often called an "*End Effector*". All the other bones assume an algorithmic position, automatically computed by the *IK* solver, to keep the chain without gaps (i.e. IK will mathematically solve the chain positions for us). This allows a much easier and precise positioning of hands and feet using IK.

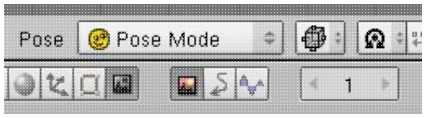
Forward Kinematics

While handling bones in Pose Mode notice that they act as rigid, inextensible bodies with spherical joints at the end. You can grab only the first bone of a chain and all the others will follow it. All subsequent bones in

the chain cannot be grabbed and moved, you can only rotate them, so that the selected bone rotates with respect to the previous bone in the chain while all the subsequent bones of the chain follow its rotation.

This procedure, called *Forward Kinematics (FK)*, is easy to follow but it makes precise location of the last bone in the chain difficult.

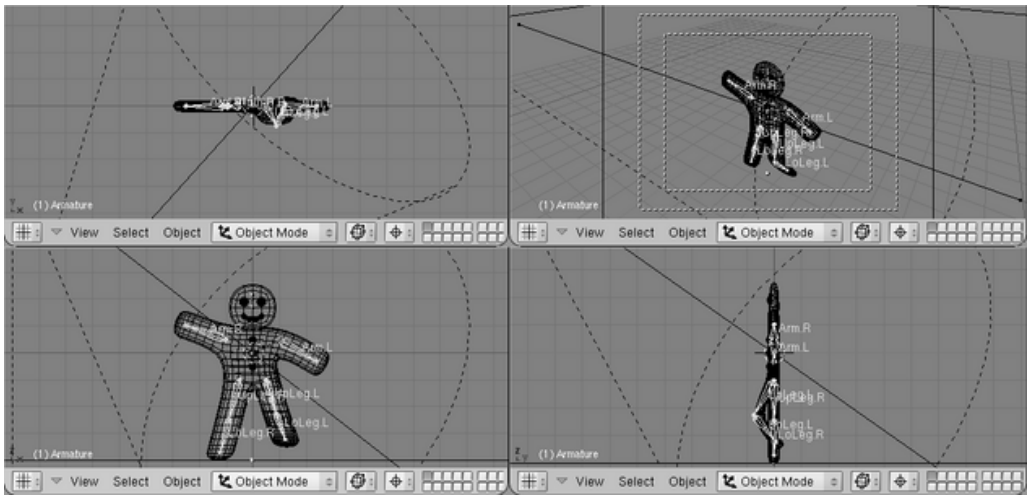
We'll make Gus walk, using FK, by defining four different poses relative to four different stages of a stride. Blender will do the work of creating a fluid animation.



The current frame Num Button in the Buttons window Toolbar.

- First, verify that you are at frame **1** of the timeline. The frame number appears in a NumButton on the right of the Buttons Window Toolbar (*The current frame Num Button in the Buttons window Toolbar*.). If it is not set to **1**, set it to **1** now.

- Now, by rotating only one bone at a time (R), we'll raise **UpLeg.L** and bend **LoLeg.L** backwards while raising **Arm.R** a little and lowering **Arm.L** a little, as shown in *Our first pose.*.



Our first pose.

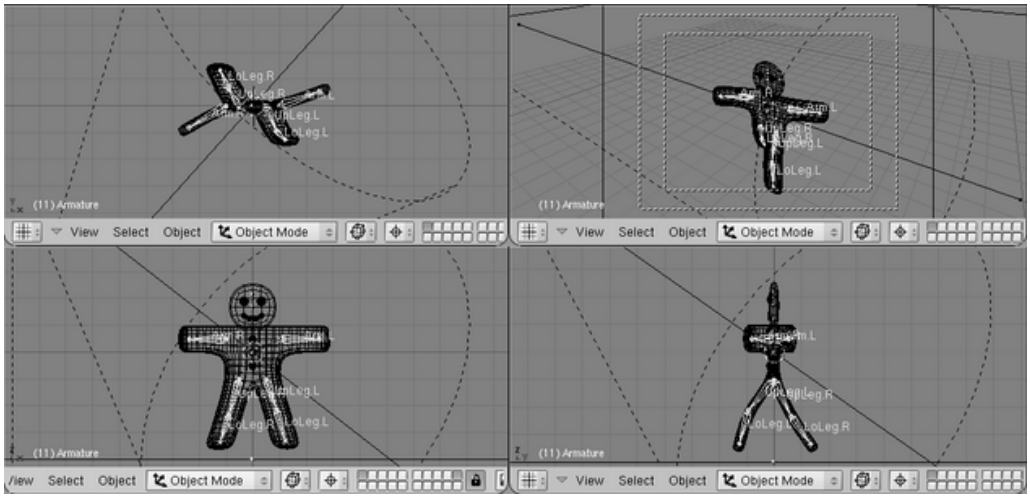


Storing the pose to the frame.

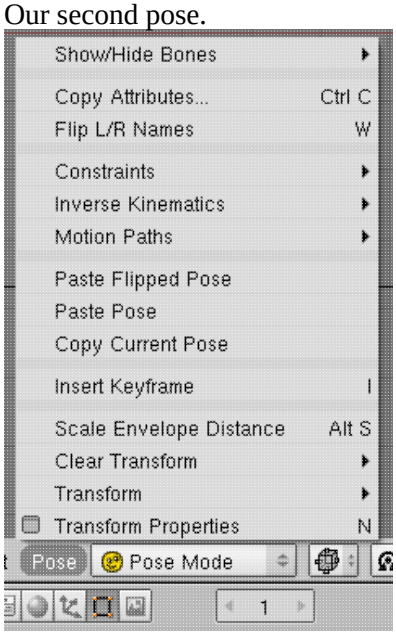
- Select all bones with A. With the mouse pointer on the 3D Window, press I. A menu pops up (*Storing the pose to the frame*). Select LocRot from this menu. This will get the position and orientation of all bones and store them as a pose at frame 1. This pose represents Gus in the middle of his stride, while moving his left leg forward and above the ground.

- Now move to frame 11 either by entering the number in the NumButton or by pressing ↑. Then move Gus to a different position, like (*Our second pose*). Start with with clearing the rotation on both arms

using AltR as mentioned earlier. From the top view, rotate **Arm.R** slightly forward and **Arm.L** slightly back. Finish the pose with his left leg forward and right leg backward, both slightly bent. Gus is walking in place!



Our second pose.



Pose menu

- Select all bones again and press I to store this pose at frame 11 and select Rot.
- We now need a third pose at frame 21, with the right leg up, because we are in the middle of the other half of the stride. This pose is the mirror of the one we defined at frame 1. Therefore, return to frame 1 and, with all the bones selected, in the Pose Menu in the 3D Window header select the Copy Current Pose entry, see (*Pose menu*). You have now copied the current pose to the buffer.
- Go to frame 21 and paste the pose with the Paste Flipped Pose option in the Pose Menu, see (*Pose menu*). This button will paste the cut pose, exchanging the positions of bones with suffix "**.L**" with those of bones with suffix "**.R**", effectively flipping it!

The pose is there but it has not been stored yet! You must press I with all bones selected.

- Now apply the same procedure to copy the pose at frame 11 to frame 31, also flipping it.

- To complete the cycle, we need to copy the pose at frame 1 *without* flipping to frame 41. Do so by copying it as usual, and by using the Paste Pose entry. End the sequence by storing the pose with I.

Checking the animation
To preview your Animation, set the current frame to 1 and press **ALT-A** in the 3D window.

Gus walks!

The single step in-place is the core of a walk, and once you have defined one there are techniques to make a character walk along a complex path. But, for the purpose of our Quick Start, this single step in-place is enough.

- Change to the Rendering Buttons (F10) and in the Anim panel, below the PLAY button, set the start frame (Sta:) to 1 (it is usually set to 1 by default so you probably won't need to change it) and set the end frame (End:) to 40 (it is set to 250 by default) (*Setting the Rendering Buttons for an animation*). Because frame 41 is identical to frame 1, we only need to render frames from 1 to 40 to produce the full cycle.



Setting the Rendering Buttons for an animation.

- Select AVI Raw as the file type in Format Panel (*Setting the Rendering Buttons for an animation.*). While this is generally not the best choice, mainly for file size issues (as will be explained later on), it is fast and it will run on any machine, so it suits our needs. (You can also select AVI Jpeg to produce a more compact file. However, it uses lossy JPEG compression and will produce a movie that some external players might not be able to play).

Finally, press the ANIM button in Anim Panel. Remember that *all* the layers that you want to use in the animation must be shown! In our case, these are layers 1 and 10.

Stopping a Rendering
If you make a mistake, like forgetting to turn layer 10 on, you can stop the rendering process with the **ESC** key.

Our scene is pretty simple, and Blender will probably render each of the 40 images in a few seconds. Watch them as they appear.

Stills
Of course you can always render each of your animation frames as a still by selecting the frame you wish to render and pressing the **RENDER** button.

Once the rendering is complete you should have a file named 0001_0040.avi in a render subdirectory of your current directory --the one containing your .blend file. The directory can be changed from the Output tab inside the Scene panel (F10).

You can play this file directly within Blender by pressing the Play button beneath the ANIM button (*Setting the Rendering Buttons for an animation*). The animation will automatically cycle. To stop it press Esc. We have produced only a very basic walk cycle. There is much more in Blender, as you'll soon discover!

Opening Files

Mode: All Modes

Hotkey: F1

Menu: File → Open

Description

Blender uses the .blend file format to save nearly everything: Objects, Scenes, Textures, and even all your user interface window settings.

Warning
Blender expects that you know what you are doing! When you load a file, you are **not** asked to save unsaved changes to the scene you were previously working on, completing the file load dialog is regarded as being enough confirmation that you didn't do this by accident. **Make sure that you save your files.**

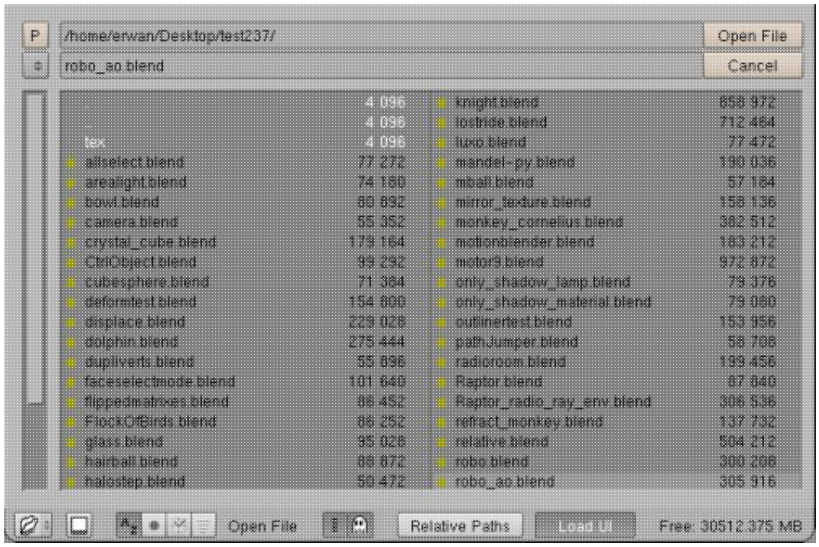
Options

To load a Blender file from disk, press F1. The window underneath the mouse pointer then temporarily becomes the File Selection window as shown in (*File Selection Window - loading.*). The bar on the left can be dragged with LMB for scrolling. To load a file, select it with LMB and then press â µ Enter, or click the Open File button. A file can also be loaded by using the MMB over the name of the file you want. It is then loaded up as if you had pressed enter.

Loading the UI

Inside each Blend file, Blender saves the user interface layout - the arrangement of screen layouts. By default, this saved UI is loaded, over-riding any user defaults or current screen layouts that you have. If you want to work on the blend file using your current defaults, start a fresh Blender, then open the file selector (F1). Turn off the "Load UI" button, and then open the file.

Navigating your Hard Disk



File Selection Window - loading.
The upper text box displays the current directory path, and the lower text box contains the selected filename. (P) moves you up to the parent directory. The button beneath, with the up and down arrow, maintains a list of recently used paths and on the windows platform a list of all drives (C:, D:, etc.). The breadcrumb files (. and ..) refer to the current directory and upper-level directory, respectively.

Other File Open Options

Open Recent
: Lists recently used files. Click on one to load it in.

Recover Last Session
: If autosave is turned on, this attempts to load the last hot backup for you. If the file cannot be found, you might need to manually go to your temp directory and find it.

Saving Files

Mode: All Modes

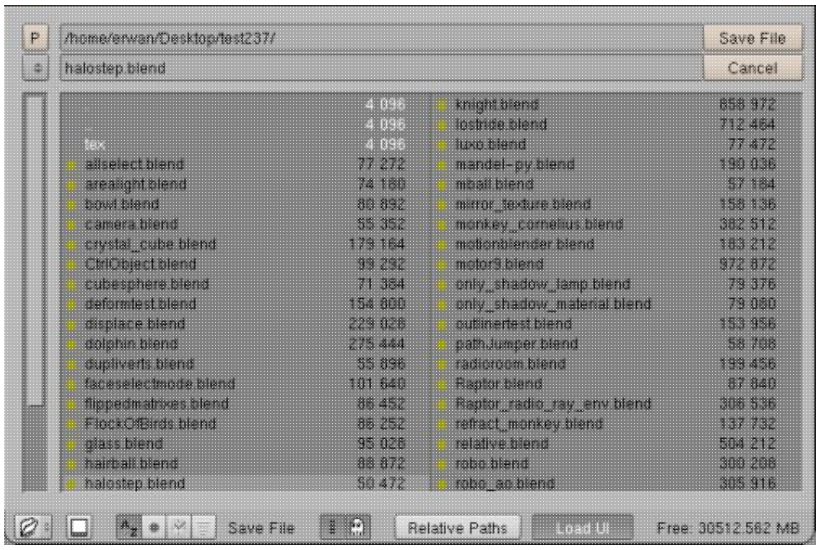
Hotkey: F2

Menu: File → Save

Description

Saving files is like loading files. When you press F2, the window underneath the mouse pointer temporarily changes into a File Selection Window, as shown in (*File Selection Window - saving.*).

Options



File Selection Window - saving.
Click the lower edit box to enter a filename. If it doesn't end with "**.blend**," the extension is automatically appended. Then press â µ Enter or click the Save File button to save the file.

If a file with the same name already exists, you will have to confirm that you want to save the file at the overwrite prompt.

Depending on the number of save versions you have set, all existing files with the same name will be rotated to a .blendX file extension, where X is 1, 2, 3 etc. So, if you were working on MyWork.blend, and saved it, the existing MyWork.blend is renamed to MyWork.blend1, and a new MyWork.blend is saved. This way, you have hot backups of old saved versions that you can open if you need to massively undo changes.

Compress Files

Enable File->Compress Files to squash large files, removing dead space.

Hints

The save dialog contains a little feature to help you to create multiple versions of your work: Pressing + NumPad or - NumPad increments or decrements a number contained in the filename. To simply save over the currently loaded file and skip the save dialog, press CtrlW instead of F2 and just confirm at the prompt.

Other File Menu Options

Append or Link
: You don't have to load a complete file; you can load in only selected parts from another file if you wish. Appending and Linking is discussed here.

Import
: Blender can use information stored in a variety of other format files which are created by other graphics programs. It does this by running a script to import the file.

Export
: Normally you save your work in a .blend file, but you can export some or all of your work to a format that can be processed by other graphics programs. To do so, you run an export script.

Rendering

Mode: All Modes

Panel: Render Context → Render

Hotkey: F12

Menu: Render → Render Current Frame

Description

This section will give you only a quick overview of what you'll need in order to render your scene. You'll find a detailed description of all options in Rendering.

Options

The render settings are in the Scene Context and Rendering Buttons Sub-context (*Rendering options in the RenderingButtons.*) which is reached by clicking the , or by pressing F10.



Rendering options in the RenderingButtons.
In the Output panel, the top field contains the path increment (default: "/tmp/") and optionally a filename prefix to use when rendering. The Path Increment is either an absolute address or a relative address. An absolute address is something like "C:\Documents\Blender\" and a relative address is a breadcrumb notation ("./" or "../") meaning to start with the current or parent directory of the Blender installation location, or a double slash ("//") meaning put the file in the directory from where the blend file was loaded.

Invalid Paths
If the construction of the path is illegal and rejected by the operating system, your file can end up in the Blender installation directory, root directory, or some other place.

The Format Panel controls the format of the render. The full size (number of pixels horizontally and vertically) and file format of the image to be created are picked here. You can set the size using the SizeX and SizeY buttons. Clicking the selection box just below the size buttons opens a menu with all available output formats for images and animations, which is currently "Jpeg" in (*Rendering options in the RenderingButtons*).

Now that the settings are complete, the scene may be rendered by hitting the RENDER button in the Render Panel or by pressing F12. Depending on the complexity of the scene, this usually takes between a few seconds and several minutes, and the progress is displayed in a separate window. If the scene contains an animation, only the current frame is rendered. (To render the whole animation, see Rendering Animations).

If you don't see anything in the rendered view, make sure your scene is constructed properly. Does it have lighting? Is the camera positioned correctly, and does it point in the right direction? Are all the layers you want to render visible? Make sure Blender Internal is chosen in the dropdown box below the RENDER

button.

Saving to disk

A rendered image is not automatically saved to disk. If you are satisfied with the rendering, you may save it by pressing F3 and using the save dialog as described in *Saving files*. The image is saved in the format you selected previously in the Format Panel.

Hints

Click the Extensions button in the Scene (F10) Render context Output panel so that Blender will add the type extension (i.e. ".jpg") automatically to image files!

Screenshots

In order to facilitate teamwork and rapid prototyping, you might want to quickly take a picture of your window or entire blender window setup.

File -> Screenshot Subwindow takes a picture of your last active window and saves it as a JPG. A window opens, allowing you to specify the location and name of the file.

File -> Screenshot All takes a picture of the entire Blender window.

Setting the default scene

If you don't like Blender's default window set-up, or want specific render settings for each project you start, or you want to save your Theme? No problem. You can use any scene file as a default when Blender starts up. Make the scene you are currently working on the default by pressing CtrlU. The scene will then be copied into a file called .B.blend in your home directory.

You can clear the working project and revert to the default scene anytime through the menu entry File>>New or by pressing CtrlX. But remember to save your changes to the previous scene first!

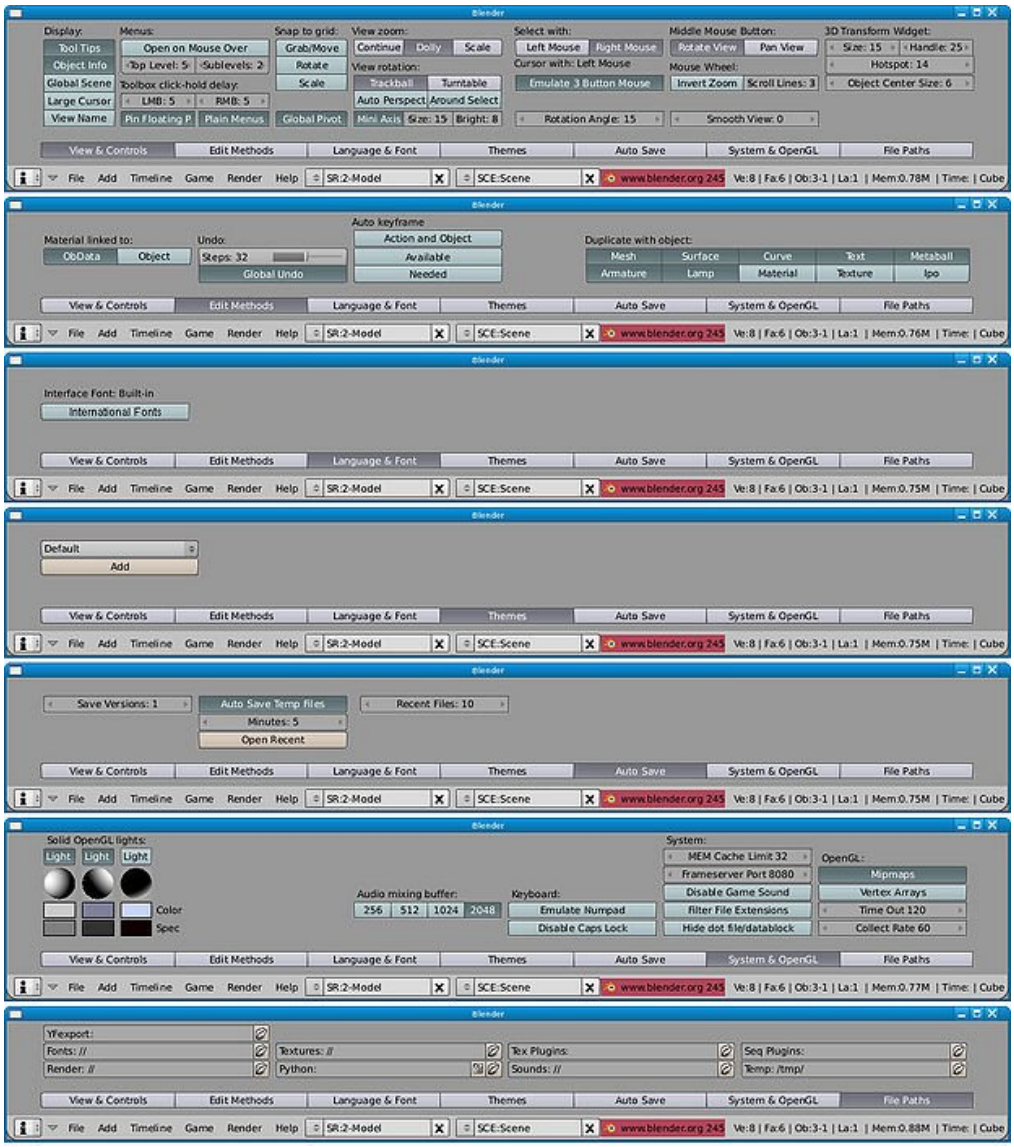
User Preferences and Themes

In Blender, you can customize your defaults, and once you are satisified, save them via File->Save user Defaults. If you ever want to start completely over, simply restore "factory" settings via File->Load Factory Settings

Blender has a few options that are not saved with each file, but which apply to all of a user's files instead. These preferences primarily concern the user interface handling details, system properties like mouse, fonts, and languages.

As the user preferences are rarely needed, they are neatly hidden behind the main menu. To make them visible, pull down the window border of the menu (usually the topmost border in the screen). The settings are grouped into seven categories which can be selected with the violet buttons shown in (*User preferences window*).

Because most buttons are self-explanatory, or display a helpful tool-tip if you hold the mouse still over them, we won't describe them in detail here. Instead, we will just give you an overview of the preference categories and some suggestions.



Pull down the User Preference Window to reveal many customization options. Each section of User Preferences Window displayed in default state.

View & Controls

> Settings concerning how the user interface should react to user input, such as which method of rotation should be used in 3D views. Here you can also activate 3-button mouse emulation if you have a two-button mouse. MMB can then be input as AltLMB.

In particular, I just want to call out the Smooth View setting used in transitioning your 3D window from one view to another (e.g. from Top view to Side view). A higher value (e.g.1000) smooths the transition from view to view, instead of jumping. A very nice effect found in other packages, and is pleasing on the eye.

Edit Methods

> Lets you specify the details for the workings of certain editing commands like duplicate. You can also change the amount of undo steps and whether the undo should work globally or locally. For more information on Undo and Redo, click here.

Add New Objects:

> Enable Switch to Edit Mode if you want the Blender to automatically go into edit mode when you add an object.
> Enable Aligned to View when new objects are added, and they will be automatically rotated to face the view you are in. Otherwise, they will be global axis-aligned.

Auto key framing for animation is also controlled from here.

- Automatic Keyframing Options
 - *Auto-Keying Enabled* Automatically sets keys after a transformation of either objects or bones, removing the need to use the I key.
 - *Add/Replace Keys* The default behavior. It Adds new keys on transformation, and, if a key already existed on the Ipo on that frame, the key is replaced with the new one.
 - *Replace Keys* Will not add new keys to Ipos. This option will only replace existing keys.
 - Other Automatic Keyframing toggles
 - *Available* Only adds keyframes to existing Ipo curves. For example, with this enabled, translation and rotation can be set in the 3D view, but only a translation key will be created if there exists only a translation Ipo, but no rotation Ipo.
 - *Needed* The default behavior for autokeying is to create a key anytime a transformation occurs, in all available channels. With the *Needed* option, values that do not change between the previous and next keys do not receive new keys. In other words, if you move an object along the Y axis only, new keys will not be set for the X or Z axes as their value does not change.
 - *Use Visual Keying* Uses the Visual Keying method for objects and bones that have certain Constraints that can affect the key values. For example, setting a key on an object with a Copy Location constraint would normally set the key for it's unconstrained location. Enabling this option causes the key to be set for the constrained location.

Language & Fonts

> International Fonts, when enabled, allow you to use English fonts as well as non Latin character fonts such as Kanji and Pharsi as the labels for Blender's buttons. When enabled, you must *Select Font* which will load the font file. Optionally, you can change the default font size in picas. Blender's default is a nice, clean sans-serif font, but you can use anything you want, even Wingdings!

Along with International Fonts, you can choose your native language (English by default), and whether Tooltips, Buttons and Toolbox text should be displayed in that language. Use Textured Fonts for better display of characters.

Themes

Blender allows the utilization of Themes to define custom interface colors and icons. You can manage themes from here, and two are built-in: Default and Rounded. Many others are available from the Internet, such as Dark Alpha and GONX. Each theme is a python script, usually ending in *_theme.py. You can import other people's themes by simply copying their [.py]thon script into your Blender/Scripts directory and restarting Blender. Once Blender has restarted, switch to the script window, click "scripts" and choose "themes", then you can select each theme that is there one at a time. Close Blender. When you restart Blender, you will see your themes listed under user preferences -> themes. Finally, press CtrlU to save it as your default.

To preserve your current theme when opening a new file, disable the Load UI button.

Customizing Theme

Click on a custom theme from the selection box, or Add a copy of the default. When you do, many more columns of buttons will be shown. The second column of buttons, in order:

> **Name** of the theme. Click into the field and change what is there)
> **Section** of the theme to change, mostly organized by window type, and has the tool tip "Specify theme for...". (3D View is the default choice)
> **Element** within that part to customize (Background is the default choice)

The next column over shows the current color for that element. Change it using the RGB sliders, or clicking on the color swatch and using the eyedropper.

Creating / Saving a Theme

File -> Export ->Save Current Theme

Customizing Icons

Blender uses LOTS of icons to represent functions; clicking an icon invokes that function. You can change many icons to suit your preference.

Default
All images in this wiki have been screenshot using the Default icon and theme, so that it all looks consistent and less confusing. If you are looking at a private tutorial, or a forum sceenshot, keep in mind they might have changed all the colors and icons, and you will have to match up based on icon placement.

To use custom icons within Blender, first create an icon file. It is a graphical image, created as a Blender Render or your favorite graphical image manipulation program as a 14x22 pixel image. Create a directory "icons" inside the ".blender" directory in Your Blender's installation. Copy the icon file set there.

In Blender, enter the "Themes" area in the "User preferences" window. Click on a custom theme, or Add a copy of the default. When you do, many more columns of buttons will be shown as described above. In the second column, the top field allows you to change the theme name. The button below it allows you to select what part of the theme to change and has the tool tip "Specify theme for...". Select 'UI and Buttons' from the drop down selection that reads '3D View'. After you did this you can select 'Icon File' from the drop down selection that's named 'Outline' (the Element to change within that . A new selection next to 'UI and Buttons' will appear and you can select your icon set there.Now You can browse icons that are in mentioned in the "icons" directory and use your favourite ones!

Auto Save

The creative process is very involving, and the artist often gets so deep into modeling and animation that he or she often forgets to bathe, eat, and especially save copies of their work. A computer crash, power outage, or simply taking a bad fork in the creative path can result in lost work or corruption of the desired product. Have no fear of immersing yourself, because Blender provides several ways to automatically save backup copies of work in progress. This sub-panel allows you to configure the two ways that Blender allows you to regress to a prior file version.

Save Versions
: The "Save Versions" button tells Blender, when you *manually* select File/Save, to save the specified number of previous versions of your file. In your current working directory, these files will be named .blend, .blend1, .blend2, etc. on up to the number of versions you specify, with the older files named with a higher number. Typically, 9 versions are more than sufficient.

Auto Save Temp Files
: Clicking the "Auto Save Temp Files" button tells Blender to *automatically* save a hot backup copy of your work-in-progress to the temp directory. Selecting this button reveals two more buttons. The first, "Minutes" button specifies the number of minutes between automatic saves. The second "Open Recent" button allows you to open the most-recent auto-save file. The auto-save file is named using a random number, has a .blend extension, and is placed in the Temp directory (refer to the "File Paths" tab). We recommend that you use these buttons to autosave into your temp folder filepath, and set how many minutes go between automatic saves (5 to 10 minutes is sufficient).

 Then, when you have done something terrible to your beautiful model, you have four choices: 1) keep working forward and try to cover up or build on your accident, 2) undo with CtrlZ, 3) regress to (open) a previously saved version in your working directory, or 4) regress to the prior auto-saved version (which is where the next button comes in). To regress to the last auto-save version, simply click the "Open Recent") button and the most recently saved work-in-progress version from the Temp directory will be loaded.

Warning
Clicking the "Open Recent" button will immediately load the most recent save, and you will lose any changes that you have made in the intervening minutes.

 Upon loading the Temp version, you may File/Save it over the current file in your work directory as a normal .blend file.

Exit No Save

When you close Blender and exit the program, Blender does not prompt you to save your work if changed. However, it automatically saves the current work-in-progress in a file called "quit.blend" in your Temp directory. If you realize that you have forgotten to save before exiting Blender, simply manually navigate to the Temp directory and open the "quit.blend" file, and save it over your work file.

Recent Files
: When you use File -> Open Recent, this control specifies how many recent files to keep track of, and how many will appear in that list.

Save Preview Images
: TBD.

System & OpenGL

Solid OpenGL lights
: When working in solid view, you can "light" the objects in your workspace with up to three light sources. The direction of each source is set by dragging inside the light sphere, and the colors of the diffuse and specular shading is set by clicking on the color swatch and using the popup color picker

Auto Run Python Scripts
: When enabled, selecting a script from any menu runs that script. However, Python is a powerful language and If you suspect a virus or get a blend file from an untrustworthy source, disable this prior to opening the blend file to prevent the script from running.

Win Codecs
: Codecs are routines that encode and decode video streams. Enable this to use any codec found in your system. Some codecs can mis-behave, and can falsely assert themselves. If you have one of these pontificating codecs on your system, disable this.

Color range for weight paint
: When weight painting, by default we use a colorful band of color to represent weights ranging from 0.0 to 1.0. If you want to use a different set of colors to represent the range of values when weight painting, click *ColorBand* and use the colorband control to set the colors that correspond to the range.

Audio Mixing Buffer
: When mixing audio or using audio in the game engine, this allows you to set aside memory for sound.

Verse
: Verse allows multiple users to work on the same blend file. Enter the URL of the Master and your username here. Then use the File->Verse menu to get connected to your buddies.

Keyboard
: if you don't have a numerical keypad and want to emulate it (for laptops). if your fat fingers keep pressing the darn caps lock instead of the tab key, disable it.

System
: Prefetch allows Blender to get anticipated frames into memory, allowing you work smoother. If you are using the Video Sequence Editor a lot, you might want to increase your MEM Cache Limit. If you are rendering frames on request as a slave, enter the port number to listen for requests on.

OpenGL
: Allow you to fine-tune how OpenGL displays textures and functions on your system.

File Paths

Choose the default paths for various file load dialogs. Remember that the // at the beginning of a pathspec means "where the .blend file is currently saved". "/" at the beginning means the top directory of the currently active drive.

Relative Paths Default
: This button is at the top right of the panel. If you enable this button, the internal path to any ancillary file, such as an image texture, will be saved, not as an absolute path starting with the drive letter, but as a pathspec to the file starting with the location of the .blend file. This enables you to zip and move entire projects from one PC to another, and all ancillary files will be found, even if on PC A you save the .blend on drive C: and on PC B you save the .blend on drive D:. Use this feature if you think you will have to be sharing or working on the project across multiple machines.

YFexport
: Blender communicates with YafRay by exporting an XML file. This entry tells Blender where to save it. Suggestion: **C:\tmp**

Fonts
: Where to find TrueType Fonts. Suggestion: **C:\Windows\Fonts**

Render
: Where to put rendered output. Suggestion: **//render**

Textures
: Where to find pictures and images for texturing surfaces. Suggest: **C:\Blender\lib\tex** where **lib** is your local library. As an alternative, you can use a local copy of the textures relative to your project. This is handy if you are going to be moving the .blend file from machine to machine and want to pack the textures. In this case, you would want to enter **//textures** to mimic the folder that is created when you unpack textures using the *write files to current directory* option.

Python
: Blender's customization scripting language and functionality extensions. If left blank, Blender uses the distributed scripts which are located in your install directory under *Blender\.blender\scripts* directory. Suggestion: **C:\Blender\scripts**. The **Python** pathspec should NOT end in a slash; this is an exception. If you use a local library/repository of scripts, you should remember to refresh it with the latest distributed scripts when you upgrade Blender.

Tex Plugins
: Plugin DLL's to augment texturing. Suggest: **C:\Blender\bin\plugins\texture**

Sounds
: wave files for soundtracks and sound effects. Suggest: **C:\Blender\lib\wav**

Seq Plugins
: When using the Video Sequence Editor within Blender, Blender has a host of nifty effects that can be augmented by DLLs. Suggest: **C:\Blender\bin\Plugins\sequence**

Temp
: The generic trash can, where your exist session save file is saved (quit.blend) as well as autosave files. Suggest: **C:\tmp**

You can manually enter a path, or LMB  click the little file folder icon to the right of the field to use Blender's file browser to navigate you hard drive/network. Doing this is recommended in order to avoid typo's. The folder searcher puts in the pathspec ending in a slash.

If any folder name changes, or the folder is moved or deleted, you will have to come back here and change them again, since Blender has no way of being informed of those changes.

Saving your Preferences

When you press CtrlU, you will save a file called .B.blend in the .blender folder underneath your Blender installation that contains the present setup, including all screens and scenes. Please note that because of its weird filename, Windows OS's may try to hide it from you. Also, it might be saved in an Application Data directory specific to your User Profile. If it is not saving your changes, be sure you have security rights set to allow changes to files in the folder; this is especially the case with Microsoft Vista OS, as it definitely does not by default allow any program to change any file in the *Program Files* directory.

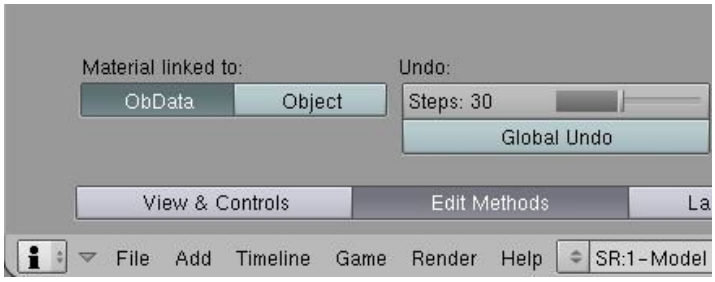
In any event, it is a plain old .blend file; so if you have objects etc in your file when you CtrlU, those will also be the default the next time you start.

If the file is lost or accidentally deleted, Blender will re-create it on the next startup.

Description

Blender has many options and features to make sure that you do not lose your work. First, it saves your actions in a list. At any time, you can tell Blender to back up in the list and undo most recent changes. Second, when you start Blender, one of the File options is to Recover Last Session. When you exit Blender, it saves the current file in a **Quit.blend** file; Recover Last Session merely loads that file back in. Third, you can tell Blender via User Preferences to automatically save versions "behind the scenes", and to keep old copies of your entire files every time you do manual saves.

Getting Started



By default, Undo is not turned off although it takes precious memory. To enable or disable undo, drag down your User Preferences window and click Edit Methods. In that panel, you may set:

- Steps: This numeric slider sets how many steps, or actions, to save. If you set this to 30, you will be able to undo the last 30 actions that you performed.
- Global Undo: This enables Blender to save actions outside of some mesh editing actions, for example, moving individual vertices while a mesh is in one editing session; each vertex move can be undone.

Undo

Mode: All Modes

Hotkey: CtrlZ

When you have done something terrible to your beautiful model, you have the following choices:

1. keep working forward and try to cover up or build on your accident,
2. Undo (via CtrlZ),
3. Revert to (open) a previously saved version in your working directory, or
4. Regress to the prior auto-saved version (if you have AutoSave turned on)

To regress to the last auto-save version, simply move your cursor toward the top of your screen until it hovers over the boundary of the User Preferences header and changes to an Up-Down arrow. Click and drag that header down to reveal the User Preferences window. Click the AutoSave tab, and click the "Open Recent" button. Immediately, the most recently saved work-in-progress version from the Temp directory will be loaded.

Warning
Clicking the "Open Recent" button will immediatly load the most recent save, and you will lose any changes that you have made in the intervening minutes.

Upon loading the AutoSaved version, you may File/Save it over the current file in your work directory as a normal .blend file.

If you have versioning turned on (again specified in the User Preferences - AutoSave tab), each time you File->Save your .blend file, Blender pushes down the previous save to a **.blend1** file, and the **.blend1** to a **.blend2** file, and so on up to the max number of versions that you have specified. To revert to one of those files, simply File->Open and select one of those **.blendx** files. Then File->Save it to bring it to the top of the stack.

Redo

Mode: All Modes

Hotkey: â § ShiftCtrlZ or CtrlY

Just as CtrlZ undoes an action, â § ShiftCtrlZ re-does the last undone action(s).

History

Mode: All Modes

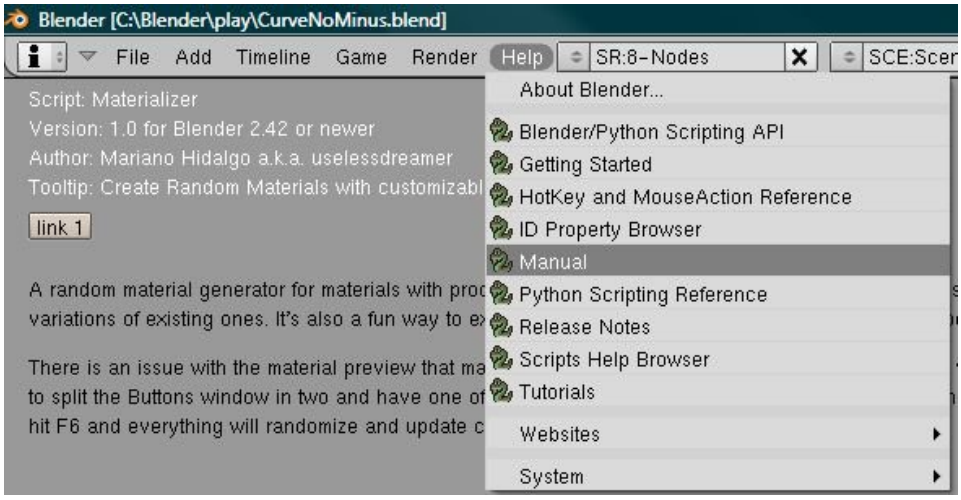
Hotkey: AltU

AltU displays the Global History of what you have done as a list of actions named generally for what you did. Clicking on any action reverts you back to that state just before the next action was performed.

Help system

Blender has an interesting mix of built-in, loadable, and web-based help, all accessed with Blender. All of the help is accessed through the Help menu at the top of the page in the User Preferences window header. Of course, any web-pages can be saved to your local hard disk or printed using your web browser for off-line viewing. We use web-based help so that we can bring you the latest up-to-date help.

Note about your web browser
Some forms of help start up your web browser and access the Blender Foundation's web servers. In order to do this, you must configure, with your operating system, a web browser as a default. If you have a dial-up connection, you must configure the web browser to automatically dial out when it starts if there is no active internet connection available.



About Blender

Displays the splash screen image, identifying the package and version.

Blender/Python Scripting API

Web-based help pages describing the application programming interface (API) that Blender exposes to Python. Python is a general programming language that can do many things on its own. In order for it to do things inside Blender, it has to access Blender functions, like creating objects, moving them, etc. Python does this by calling on specific Blender API calls, like ob.setLocation to change an object's location.

Getting Started

Web-based help pages listing specific links to official sites maintained with fool-proof, easy tutorials that can help you get started and comfortable using Blender.

HotKey and MouseAction Reference

Internal There's an old adage among Blender users:

"One hand on the mouse, the other on the keyboard."

Blender makes extensive use of hotkeys to allow you to quickly perform common actions, like R to Rotate the selected object. This internal page tells you what all those keys do, separated into Arrows, Letters, Mouse actions, Numbers, etc. It has a Search Function, so if you vaguely know the term for *what* you want to do, you can search for the hotkey to find out *how* to do it.

ID Property Browser

Internal For all the various kinds of objects in your .blend file, this screen allows you to find them (starting at the Scene level) and drill down to inspect their properties.

Manual

Web-based Brings you to the main table of contents page of the wiki.

Python Scripting Reference

Same as Blender/Python Scripting API

Release Notes

Web-based A page listing all the Release Notes, telling you what has changed and new features added to Blender, back several versions. The release notes provide you with a quick explanation of "what's new", frequently along with examples and demonstrations and movies/images.

Scripts Help Browser

Internal This window allows you to choose a script (from your \scripts directory) by category, and shows you the help information embedded within the script. Use this to find out more about what these custom Python scripts do, and how to use them.

Tutorials

'*Web-based* Lists simple, step-by-step tutorials that can provide you with easy-to-follow directions on how to use Blender. Please also visit the wiki Tutorials page.

Websites

Web-based Lists frequently accessed websites related to Blender and its development.

System

With a 3D View window active, clicking Benchmark enables you to benchmark three different actions in Blender. The resulting statistics, particularly the 'operations per second (ops/s) is useful for comparing the performance of Blender across different machines. For a more robust benchmark, refer to the [official Benchmark site.]

System Information creates an info.txt file which lists various key properties of your system and Blender, useful in diagnosing problems. These are the instructions to get it:

- run it by selecting it in Help Â» System Â» System information
- after a while a popup appears to confirm it has finished
- open the text editor
- select "system-info.txt" in the text menu selector. In case you see an apparently blank page, please try to scroll the page with the Wheel , the text cursor might be at the end of the text.

Introduction

Using Blender, you create a world that exists in four dimensions:

1. Left-right, commonly called the â← *x*â← axis.
2. Forward-backward, commonly called the â← *y*â← axis.
3. Up-down, commonly called the â← *z*â← axis.
4. Time-sensitive, through animated objects, materials, and motion captured in frames.

The problem is, you have a two-dimensional computer screen in front of you! Your mouse can only move left-right and up-down. You cannot go back in time, and you canâ← t literally reach out into the screen and grab an object and move it somewhere else (although I did see a scary movie once where the guy reached through a mirror, but thatâ← s another story).

Instead, you have to tell Blender to do it for you. This section tells you how to navigate around in your virtual world using the unique Blender interface. Have a nice trip !

Introduction

Blender lets you work in three-dimensional space, on monitor screens that are only two-dimensional. To be able to work in three dimensions, you must be able to change your viewpoint as well as the viewing direction of the scene. This is possible in all of the 3D Viewports. While we will describe the 3D View window, most other windows use an equivalent series of functions. For example, it is possible to translate and zoom a Buttons window and its panels.



Mouse Buttons and Numpad:

If you have a mouse with less than three buttons or a keyboard without numpad, please refer to the Keyboard and Mouse page of the manual to learn how to use Blender with them.

Rotating the View

Mode: All modes

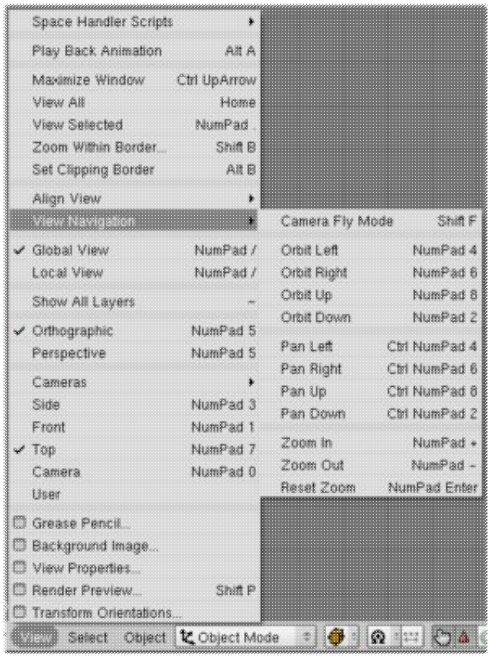
Hotkey: MMB / 2 NumPad / 4 NumPad / 6 NumPad / 8 NumPad / CtrlAltWheel

Menu: View â View Navigation

Description

Blender provides four default viewing directions: Side, Front, Top and Camera view. Blender uses a right-angled â← Cartesianâ← coordinate system with the Z axis pointing upwards. â← Sideâ← corresponds to looking along the X axis, in the negative direction, â← frontâ← along the Y axis, and â← topâ← along the Z axis. The Camera view shows the current scene as seen from the camera view point.

Options



A 3D viewportâ← s View menu.

You can select the viewing direction for a 3D viewport with the View menu entries (*A 3D viewportâ← s View menu*), or by pressing the hotkeys 3 NumPad for â← sideâ← , 1 NumPad for â← frontâ← , 7 NumPad for â← topâ← . You can select the opposite directions if you hold Ctrl while using the same numpad shortcuts. Finally 0 NumPad gives access to the â← cameraâ← viewpoint.

Hotkeys

Remember that most hotkeys affect **the active window**, that has focus, so check that the mouse cursor is in the area you want to work in before your use the hotkeys!

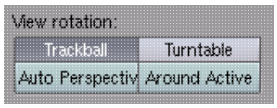
Apart from these four default directions, the view can be rotated to any angle you wish. Click and drag MMB on the viewportâ← s area: if you start in the middle of the window and move up and down or left and right, the view is rotated around the middle of the window. Alternatively, you can press and hold Alt while dragging LMB in the viewportâ← s area.

To change the viewing angle in discrete steps, use 8 NumPad and 2 NumPad (which correspond to vertical MMB dragging, from any viewpoint), or use 4 NumPad and 6 NumPad (or CtrlAltWheel) to rotate the scene around the Z global axis, *whatever being the point of view*.

TrackBall/Turntable

By default, when you rotate the view as described in the preceding section, you are rotating the scene as though you are rolling your hand across a â← **trackball**â← . For some users this is intuitive and for others it is not. If you feel you are having difficulties with this style of 3D window rotation you can switch to the â← **turntable**â← style.

The Turntable style is fashioned more like a record player where you have two axes of rotation available, and the world seems to have a better definition of what is â← Upâ← and â← Downâ← in it. The downside to using the Turntable style is that you lose some flexibility when working with your objects. However, you gain the sense of â← Upâ← and â← Downâ← which can help if you are feeling disoriented. Of course you can always switch between the styles depending on what you are working on.



View rotation.
To change the rotation â← styleâ← , use the User Preferences window (remember to pull the main window down because only the header shows by default). Click on the View & Control button to reveal a page of buttons relating to *views* and *control* functionalities. You will see an area for choosing the View rotation:, see (*View rotation*). There are two additional buttons for controlling the display in the 3D window. Auto Perspective will automatically switch to perspective whenever the view is rotated using MMB . Around Selection will rotate the view around the center of the current selection. If there is no selection at that moment (e.g. if you used A to deselect everything), the last selection will be used anyway.

Panning the View

Mode: All modes

Hotkey: â § ShiftMMB / Ctrl2 NumPad / Ctrl4 NumPad / Ctrl6 NumPad / Ctrl8 NumPad / â § ShiftAltLMB

Menu: View â View Navigation

Description

To pan the view, hold down â § Shift and drag MMB in the 3D Viewport. For discrete steps, use the hotkeys Ctrl8 NumPad, Ctrl2 NumPad, Ctrl4 NumPad and Ctrl6 NumPad as with rotating (note: you can replace Ctrl by â § Shift). For those without a middle mouse button, you can hold â § Shift Alt while dragging with LMB.



MMB default behavior.
The behavior of the MMB can be customized, in the User Preferences window, View & Control tab, so it will pan by default and â § ShiftMMB will rotate the view.

Zooming the View

Mode: All modes

Hotkey: CtrlMMB / Wheel / + NumPad / - NumPad / CtrlAltLMB

Menu: View â View Navigation

Description

You can zoom in and out by holding down Ctrl and dragging MMB. The hotkeys are + NumPad and - NumPad. The View â Viewport Navigation sub-menu holds these functions too; see (*A 3D viewportâ ← s View menu*).

If you have a wheel mouse, you can perform all of the actions that you would do with + NumPad and - NumPad by rotating the wheel (Wheel). Since the Buttons window has so many panels, rotating the mouse wheel pans the window left and right in horizontal view. This allows you to pan to the panel you need in a narrow or smaller display, or if the window is narrow. As an alternative, you can easily display the Buttons window vertically, its panels will be arranged top to bottom. If you have neither a wheel mouse nor a middle mouse button, you can easily zoom in and out with CtrlAltLMB.

If You Get Lostâ ←¦
If you get lost in 3D space, which is not uncommon, two hotkeys will help you: â Home changes the view so that you can see all objects (View â View All menu entry), while . NumPad zooms the view to the currently selected objects (View â View Selected menu entry).

Fly mode

Mode: Object mode

Hotkey: â § ShiftF

Menu: View â View Navigation â Camera Fly mode

Description

When you have to place the view, normally you do as described above.

There are cases in which you really prefer to just navigate your model, especially if it's very large, like environments or some architectural model. In these cases viewing the model in perspective mode has limitations, for example after zooming a lot panning is extremely uncomfortable and difficult, or you apparently cannot move the camera any nearer. As an example, try to navigate to a very distant object in the view with traditional methods (explained above) and see what you can get.

With Flymode you move, pan and tilt, and dolly the camera around without all of those limitations.

Options

To enter in Flymode, in Object mode, press â § ShiftF: the mouse pointer will move at the center of the view, and a squared marker will appear. Some of the options of Fly mode are influenced by the position of the mouse pointer relatively to the center of the view itself, and the squared marker at the center of the view provides a sort of "safe region" where you can place the mouse with no effect on the view. The more you take the mouse pointer away from the marker, the more you pan, or track, etc.
You can:

- move the mouse left/right to pan the view left/right or move it up/down to tilt the view up/down
- move the view forward/backward:
 - ♦ WheelUp or + NumPad to accelerate the view
 - ♦ WheelDown or to - NumPad to decelerate the view

 (if the view is already moving forward, decelerating can stop it or even move it backward)

- drag the MMB to dolly: in this case the view can move laterally on its local axis at the moment you drag the mouse: quite obviously, dragging left/right/up/down makes the view dolly on the left/right/up/down respectively.

When you are happy with the new view, just click LMB to confirm. In case you want to go back from where you started, just press Esc or RMB.

As an interesting note, you can aim the camera in fly mode too.

The player will show in this paragraph

Perspective and Orthographic Projection

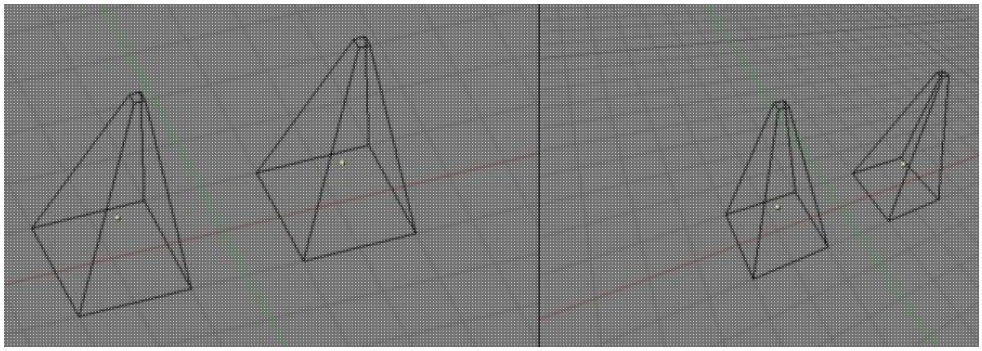
Mode: All modes

Hotkey: 5 NumPad

Menu: View â Perspective/Orthographic

Description

Each 3D Viewport supports two different types of projection. These are demonstrated in (*Orthographic (left) and perspective (right) projection.*):

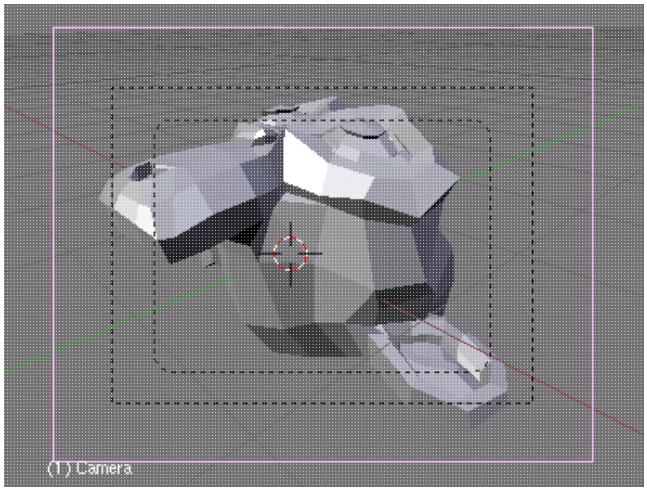


Orthographic (left) and perspective (right) projection.
Our eye is used to perspective viewing because distant objects appear smaller. Orthographic projection often seems a bit odd at first, because objects stay the same size independent of their distance. It is like viewing the scene from an infinitely distant point. Nevertheless, orthographic viewing is very useful (it is the default in Blender and most other 3D applications), because it provides a more â← technicalâ← insight into the scene, making it easier to draw and judge proportions.

Options

To change the projection for a 3D Viewport, choose the View â Orthographic or the View â Perspective menu entry. The hotkey 5 NumPad toggles between the two modes.

Camera projection
Changing the projection for a 3D Viewport does not affect the way the scene will be rendered. Rendering is in perspective by default. If you need to create an orthographic rendering, select the camera and press Orthographic in the Camera panel, Editing context (F9).



Demonstration of camera view.
The View â Camera menu entry sets the 3D Viewport to camera mode (Hotkey: 0 NumPad). The scene is then displayed as it will be rendered later (see *Demonstration of camera view*): the rendered image will contain everything within the outer dotted line. Zooming in and out is possible in this view, but to change the viewpoint, you have to move or rotate the camera.

Technical Details

Perspective definition

A *perspective* view is geometrically constructed this way: you have a scene in 3D and you are an observer placed at a point O. The 2D perspective scene is built by placing a plane, a sheet of paper where the 2D scene is to be drawn in front of point O, perpendicular to the viewing direction. For each point P in the 3D scene a line is drawn, passing from O and P. The intersection point S between this line and the plane is the perspective projection of that point. By projecting all points P of the scene you get a perspective view.

Orthographic definition

In an *orthographic* projection, you have a viewing direction but not a viewing point O. The line is then drawn through point P so that it is parallel to the viewing direction. The intersections S between the line and the plane is the orthographic projection. And by projecting all points P of the scene you get the orthographic view.

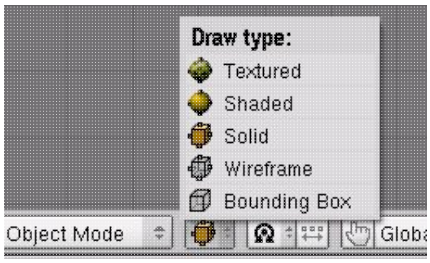
View Shading

Mode: All modes

Hotkey: Z / â § ShiftZ / AltZ / D

Description

Depending on the speed of your computer, the complexity of your scene, and the type of work you are currently doing, you can switch between several drawing modes:



A 3D Viewportâ ← s draw mode button.

Textured (AltZ)
: Displays UV image textured models with OpenGL lighting. Procedural textures will not be shown.

Shaded (â § ShiftZ)
: Approximates all textures and lighting at each vertex, and blends from one to the next. Much less accurate than using the render engine to check textures, but much faster. Note that if you have no lighting in your scene, everything will remain black.

Solid (Z or AltZ)
: Surfaces are drawn as solid colours, with built-in OpenGL lighting (not dependent on scene light sources).

Wireframe (Z or â § ShiftZ)
: Objects only consist of lines that make their shapes recognizable. This is the default drawing mode.

Bounding Box
: Objects arenâ ← t drawn at all; instead, this mode shows only the rectangular boxes that correspond to each objectâ ← s size and shape.

You can also pop up a contextual menu by pressing D to select between all the draw modes.

Local View

Mode: All modes

Hotkey: / NumPad

Menu: View â Local View

Description

When in local view, only the selected objects are displayed, which can make editing easier in complex scenes. To enter local view, first select the objects you want (see Selecting objects), and then use the View â Local View menu entry. Use the View â Global View menu entry to go back to Global View. The hotkey / NumPad toggles between Local and Global View.

View Clipping Border

Mode: Any mode

Hotkey: AltB

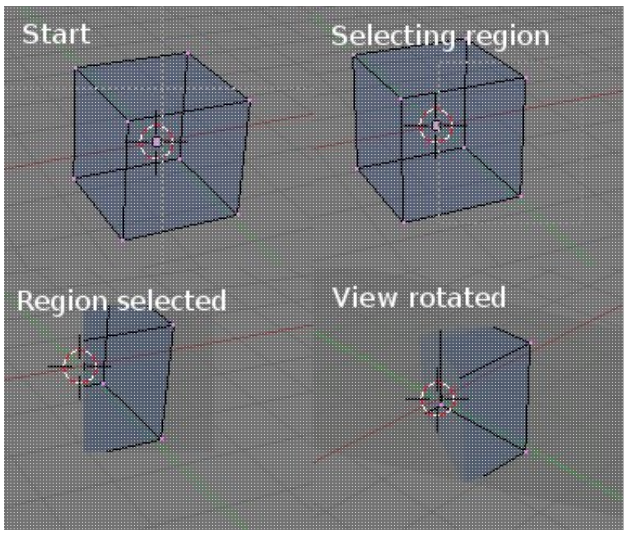
Menu: View â Set Clipping Border

Description

To assist in the process of working with complex models and scenes, you can change the view clipping to visually isolate what youâ← re working on.

Using this tool you create a 3D clipping volume shaped like a pyramid; you could think of it as a Frustum with a top. You specify the base of the pyramid by creating a 2D rectangular region.

Examples



Region/Volume clipping.
(*Region/Volume clipping*) is an example of using the clipping tool with a cube. Start by activating the tool with AltB, see â← Startâ← in the upper left. This will generate a dashed cross-hair cursor. Click with the LMB and drag out a rectangular region shown in the upper right. Now a region is defined and clipping is applied against that region in 3D space. Notice that part of the Cube is now invisible or clipped. Use the MMB to rotate the view and you will see that only what is inside the pyramid volume is visible. All Edge/Face tools still function as normal but only within the pyramid volume.

The gray area surrounding the volume is the pyramid volume itself. To deactivate the clipping tool toggle it by applying AltB a second time. All of 3D space will become visible once again.

Introduction

The 3D View is where you perform most of the object modeling and scene creation. Blender has a wide array of tools and options to support you in efficiently working with your mouse, keyboard and keypad. Your flat (two-dimensional) monitor is your viewport into the 3D space.

It is also the oldest, and therefore most feature- and option-rich areas of Blender. **Donâ← t be intimidated!** Most of us humans never use all the features here, just like we donâ← t use all of our diction on a daily basis either. So, take it slow, a few at a time, experimenting to see what they do.

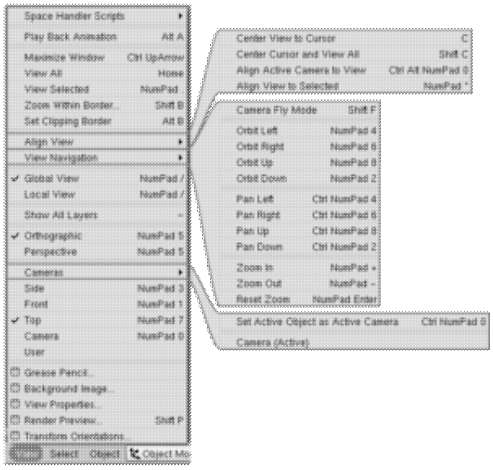
3D Window Header

The 3D View window is comprised of a workspace and a header. The header is shown at the top or bottom of the workspace, and can be hidden if desired. The header shows you a menu and the current mode, as explained below.



3D View header.

View Menu



The View menu.

This menu provides options to control the way the workspace is viewed:

- Space Handler Scripts - This submenu shows Space Handler Scripts; by default, there arenâ ← t any.
- Play Back Animation - This item plays back the animation from the current frame.
- Maximize Window - This item maximizes the 3D View window to fill the entire Blender window, and once selected this menu item will change to Tile Window, if menu entry Tile Window is then selected the 3D View window will be restored to itâ ← s previous size. Note that this feature is available for nearly every window types in Blender. As well as the menu entry, the shortcuts Ctrl↑ and Ctrl↓ can also act as toggles to maximize/tile not only the 3D View window, but also any window which currently has focus. In addition to these shortcuts, it is also possible to maximize/tile the currently focused window with shortcut â § ShiftSpace, making it extremely convenient for laptop users, as they can quickly maximize/tile the currently focused window to work on another window such as the Buttons/Outliner windows, for example.
- View All - This command zooms the 3D view to encompass all the objects in the current scene.
- View Selected - This command zooms the 3D view to encompass all the selected objects.
- Zoom Within Border... - This command allows you define the area you want to zoom into.
- Set Clipping Border - This command allows you to define a clipping border to limit the 3D view

display to a portion of 3D space. For more information on this command, see View Clipping Border in this chapter.

- Align View - This submenu allows you to shifts your view to be centered on the cursor (Center View to Cursor, C). Center Cursor and View All (â § ShiftC) centers your view **and** zooms out so that you can see everything in your scene. You can also change your viewpoint to be through the active camera, and center the camera on the cursor (Align Active Camera to View, CtrlAlt0 NumPad). Instead of the cursor, you can center your view on the selected object (Align View to Selected, * NumPad).
- View Navigation - This submenu contains commands for rotating and panning the view. Using these commands through the menu is not that efficient; however, like all Blender menus, the much-more-convenient keyboard shortcuts are listed next to the commands. Camera Fly mode moves your view through 3D space. Use the keys indicated to orbit your view, or hold down MMB and move your mouse. Use the keypad + NumPad/- NumPad keys to zoom, or scroll your mousewheel. See this page for more about navigating in 3D views.
- Global View/Local View - Global view shows all of the 3D objects in the scene. Local view only displays the selected objects. To switch between global and local view use / NumPad. Accidentally pressing / NumPad can happen rather often if youâ ← re new to Blender, so if a bunch of the objects in your scene seem to have mysteriously vanished, try turning off local view.
- Orthographic, Perspective - These commands change the projection of the 3D view. For more information, see Perspective and Orthographic Projection. Generally you want to stay in Orthographic view.
- Show All Layers - This enables all of the twenty layers in the 3D view, allowing you to see **all** objects that are in your scene (especially if you combine it with a View All action). When enabled, this item turns to Show Previous Layers, to go back to previous layers selection.
- Cameras - This submenu lists all the cameras in the scene. Selecting one will make it the active camera. There is also a command (Set Active Object as Active Camera, Ctrl0 NumPad) that sets the current object (which doesnâ ← t have to be a camera) as the camera, so you can see what the scene looks like from its point of view.
- Side, Front, Top - These commands change the view to the default side, front, or top views. Pressing the Ctrl key changes to the â ← complementaryâ ← view: Ctrl3 NumPad for right side, Ctrl1 NumPad for back, or Ctrl7 NumPad for bottom-looking-up views.
- Camera - This command switches the view to the current camera view.
- User - This command switches to a user view. In most cases, this wonâ ← t seem to do anything, but if you are in the camera view or have orthographic projection on, the view will change to perspective (and leave the camera view, if applicable).
- Grease Pencil... - This command will toggle the Grease Pencil floating panel, see this page.
- Background Image... - This command will toggle the Background Image floating panel, which allows you to load and pick an image to display in the background of the orthographic 3D view, as well as adjust its size and position. This is useful if you have a picture (for example, a face) that you want to model from. Each pane (3D View window) has its own background image settings. Each pane can use (or not use) a background image independently. So, you can set Top view to have one image, Side view can have another, and Front another. Each view uses an image separately from the others. However, all can use be the same image if the image is one big composite of all views you want to reference; just use the offset values in each pane to position the image where you want it. Background images can be stills, movies (avi or sequences) or even a render from another scene. For movies, enable Auto Refresh and Blender will display the appropriate frame from the movie when you change frames in your animation (with an optional offset, Offs field).



Use Lo-Res Proxy:

To improve PC performance when using background images you may have to use lower-resolution proxies. If your monitor resolution is 800Ã 600, then the background image, full screen, without zooming, only needs to be 800Ã 600. If your reference image is 2048Ã 2048, then your computer is grinding away throwing away pixels. Try instead to take that 2kÃ 2k image, and scale it down (using Blender, or Gimp) to, for example, 512Ã 512. You will have sixteen times the performance, with no

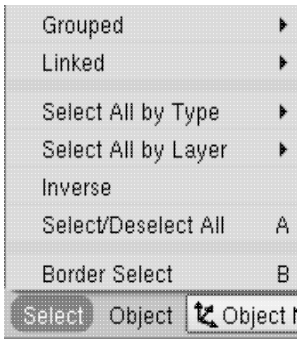
appreciable loss of quality or exactness. Then, as you refine your model, you can increase the resolution.

- View Properties... - This command toggles the View Properties floating panel, which allows you to toggle the grid and adjust its spacing, adjust the zoom of the camera, toggle specific axes (X, Y, and Z), view and change the specific location of the 3D cursor, adjust several toggles (outlining the selected object, showing all objectsâ← centers, showing relationship lines), and, lock the 3D view so that it always points towards a certain object or bone. See the View Properties window references for more details.
- Render Preview... - This command toggles the Preview panel, which shows a (relatively) live preview render of whatever it is over.
- Transform Orientation... - This command toggles the Transform Orientations floating panel, which allows you to select the coordinate system you want to use for your transformations. You can do the same with Orientation drop-down menu in the header, but this panel has an extra option, which might be *very* useful in some situations: you can define other coordinate systems, based on the local systems of selected objects. See also this page.

- Onion Skinning/Ghosting - Selecting keyframe display mode (Show Keys, K) in the Ipo window, and possibly then pressing K again in the 3D View (depending on your PCâ← s OpenGL support), will display the keyframed locations for an object (mesh, armature, etc.). Showing the ghost past/present keyed positions Ipo Keys is an often overlooked feature that can greatly assist animation visualization. The location of the object at the current frame is shown as a green line in the Ipo window, and as the object in the 3D View. The keyframe selected in the Ipo window is shown in yellow, as is the outline of the object in the 3D View, further helping you visualize the animation. All other ghost keyframed locations are shown as a black outline in the 3D View. Ghosting for armatures is more versatile and thus more complicated; see the Armature Visualization options.

In addition, the * NumPad key orients the view normal to the selected face (in Edit modeâ←¦).

Select Menu



The Select menu.

- Grouped - Blender has a few ways of grouping items together. This submenu contains commands to select objects by their various groupings.
 - ♦ Children - This command selects **all** the children of the selected object(s) (i.e. their children, the children of their children, etc.).
 - ♦ Immediate Children - This command selects the children of the select object(s); however, unlike the previous command, it does not continue selecting the children of the children, just the direct descendants of the parent.
 - ♦ Parent - This command selects the parent(s) of the selected object(s).
 - ♦ Siblings (Shared Parent) - This command selects all the objects that share the parent of an

object. This means that if you have, for example, several Blender objects that make up one physical object and are all children of one part of this object (like an empty), you can select one Blender object that is part of the physical object, use this command, and all the Blender objects that make up the physical object will be selected (if that makes any sense!).

- ♦ Objects of Same Type - This command selects all objects of the same type (Lamp, Mesh, Camera, etc.).
- ♦ Objects on Shared Layers - This command selects all the objects on the same layer(s) as the selected object(s).
- ♦ Objects in Same Group - This command selects all the objects in the groups that the selected object(s) are in. This can be used for the same purpose as Siblings (Shared Parent) if you have grouped parts of a physical object instead of using parents.
- ♦ Objects Hooks - This command selects any objects that are acting as Hooks for the selected object(s).
- ♦ Object PassIndex - This command selects all the objects that use the same pass index (setting in the Object context â Object sub-context â Object and Links panel â PassIndex numbutton. See also the Render Passes and Convertor Compositing Nodes pages).
- ♦ Object Color - This command selects all the objects that use the same â← object colorâ← (setting in the Object context â Object sub-context â Draw panel â color picker, upper right corner).
- ♦ Game Properties - This command selects all the objects that have game properties of the same name as those of the selected object(s).

- Linked - This submenu contains commands that allow you to select objects based on datablocks (Ipo curves, Materials, Textures, etc.) that they share.
 - ♦ Object Ipo - This command selects all the objects that share the Object Ipo curves of the selected object(s).
 - ♦ ObData - This command selects all the objects that share the ObData datablocks of the selected object(s) (this means, the same meshes, curves, â←¦).
 - ♦ Material - This command selects all the object(s) that share the material(s) of the selected object(s).
 - ♦ Texture - This command selects all the object(s) that share the texture(s) of the selected object(s).
 - ♦ If you use the shortcut (â § ShiftL) instead of the menu, you will have two more options: DupliGroup will select all the objects that duplicate the same group(s) as the selected one(s), and ParticleSystem will select all the objects that use the same particle system(s) as the selected one(s).
- Select All By Type - This submenu contains commands that allow you to select all the objects of a certain type (Mesh, Camera, Lamp, etc.)
- Select All By Layer - This submenu contains commands that allow you to select all the objects in a specified layer.
- Random - As you might have guessed, this command randomly selects some objects in the scene. The percentage you have to specify is not the amount of objects that will be selected, but the probability each object has to be selected (subtle, isnâ← t it? For example, if you specify 50%, this does not mean that half of your objects will be selected, but that each object has a fifty-fifty chance to be selected: you might have only one object, or nearly all objects selected, even though this is not very probableâ←¦).
- Inverse - This command inverts the selection (selects all the deselected objects and deselects all the selected objects).
- Select/Deselect All - This command deselects the current selection if there is one; if nothing is selected, it selects everything.
- Border Select - This command allows you to select objects using the traditional rectangle that most programs use. After pressing B or selecting this menu option, click LMB and drag your mouse diagonally through your workspace. When you release your LMB, all objects that were partly or totally within the box will be selected. If you instead click MMB or RMB and drag, all objects within the box will be *de*-selected.

Object Menu



The Object menu.
This menu operates on objects as a whole. Many options act on the active object, based on other objects. You indicate that the Blender by RMB clicking on the base object, then â § ShiftRMB clicking on the active object, and then invoking the menu option. In the case of wanting to work on more than two objects, simply click on all the base objects first; the last object selected will be the active object.

- Scripts - This submenu lists available Python scripts written to extend blenderâ← s object handling capability. Of note is the Knife tool that cuts meshes. See each scriptâ← s documentation for more information on how to use the script.
- Show/Hide Objects - This submenu allows you to Show Hidden objects (AltH), to Hide Selected objects (H), or to Hide Deselected objects (â § ShiftH).
- Move to Layer... - Objects can be organized into 20 Layers, and only a few layers can be selected for display to avoid clutter. This option moves the selected object(s) to a layer. To do so, select one of the 20 layers by clicking on one of the buttons in the popup window and click the OK button. If the layer is not selected for display, the object is removed from view. To view the objects on a layer, click the appropriate layer button (located to the right on the header).
- Convert Object Type... - This command pops-up a menu, different for each type of objects. For example, Curve and Surface objects can be converted to Mesh, and you can apply the modifiers of a Mesh object, keeping or not the original object.
- Join Objects - Joins multiple selected objects into on single object.
- Boolean Operation... - This command, only available with mesh objects, pops-up a menu that allows

you to perform discrete operations on the active object based on previously selected objects. Union (OR) extends the object to include all selected objects. Difference (XOR) modifies the active object by cutting away parts where they intersect. Intersect (AND) discards everything except where they intersect. The first three options are destructive. If you use modifiers, they can be undone or applied at any moment.

- Constraints - You can constrain an object based on another object, like a dog on a leash. Various ways to constrain the active object (based on another previously selected object) are by location, rotation, scale, etc.
- Track - You can make an object face, or always point to another object by the Track option. If you move the base object, the active object will rotate so that it always â← keeps an eyeâ← on the base object (note that you should use the new â← trackâ← constraints, not the old system).
- Group - Grouping is a completely arbitrary way for you to group like things together. If you are making a scene of a park, you can group all the trees together if you like.
- Parent - Use this submenu to designate the active object as the parent of one or more child objects. Then, when you move the parent, the children move with it.
- Copy Attributes - An object is in a certain place, called a location. â← Locationâ← is an attribute of the object, as well as its size, rotation, constraintsâ←¦ All selected objects can copy the attributes of active object using this option.
- Make Local - If an object was linked from another scene, this makes a local copy of it for the current scene.
- Make Single User - Make a â← localâ← copy of some of the datablocks shared by the object with others. For example, an 8-ball object can share the same mesh (called â← `billiards ball`â←) with 15 other users (15 balls plus the cue ball). Making a local copy assigns a copy (by default suffixed with the first number available, like â← `billiards ball.001`â←) of the multi-user mesh to the selected object. You can then edit the local copy without affecting any of the other users.
- Make Links â←¦
 - To Scene... - This is a way to proxy an object to another scene. It exists in the current scene, but will show up (**and be modifiable**) in the linked scene as well.
 - Object Ipo, Mesh data, etc. - It is a way to make the current object join the multi-user community and share the selected item with the other selected objects (**be careful**, it means that you assign, for example, the material of the active object to all selected objects).
- Delete - Wipes it out, man. Like totally. If its material and texture were single user and now are not used, they will not be saved in the .blend file.
- Duplicate Linked - Makes a copy of the object, but links their mesh as multi-user, so if you change one (like a table leg), all the other table legs match.
- Duplicate - Makes a xerox copy. Well, a three-dimensional xerox copy, if there is such a thing.
- Insert Keyframe - Records in Ipos the current location, rotation, etc (whatever you select) of the object at that frame. Use this to set up basic animations.
- Snap - This menu allows you to move the selected object to the cursor, grid or vice versa. Very handy in making items share the same space.
- Clear/Apply - Clears (resets to zero) the objectâ← s scale/rotation as selected, or applies current rotation/scale to the object making them default.
- Mirror - Applies a mirror transformation on the object.
- Transform - Use the menu to refresh your memory of the most common hotkeys.
- Transform Properties - Pressing N pops up a floating panel that gives you key information about the active object: location, rotation, scale, dimensionâ←¦

Mode List

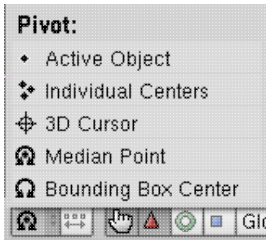


The Mode drop-down list.
Blender has a few modes of operation. Working on Objects as a whole, or in Edit mode by allowing you to modify the shape of the object. In Sculpt mode, your cursor becomes a tool to shape the object, while your cursor becomes a brush in Vertex Paint, Weight Paint, or Texture Paint modes.

ViewPort Shading List

See the previous page.

Rotation Selector



Pivot point selector.
When rotating or scaling an object or bunch of vertices/edges/faces, you may want to shift the pivot point in 3D space. Using this selector, you can change the pivot point to the location of the cursor, the average center spot of the selected items (median), etc. Pivoting is fully described here.

Use the â← Object Centerâ← toggle to switch between multi-point averages or last selected as the center.

Transform (Manipulator) Selectors

These handy selectors, also featured in other not-to-be-named CG packages, allow you to rotate or move objects by grabbing (clicking with your mouse) their controls and moving your mouse in the axis. Each color stands for an axis.

Click on the little â← fingerâ← icon to enable the manipulator. Your object in 3D View will now have the manipulator around its center. The **transform** manipulator is a 3-axis arrow; the **rotate** manipulator are three circles, one for each axis; click and drag on the arc to rotate on an axis. The **scale** manipulator is 3-axis of lines that end in a block. Customize the size of the manipulators in the User Preferences.

LMB click the buttons for a 3D Move, Rotate, or Scale Selector. â § ShiftLMB to activate multiple manipulators at the same time. You can move, rotate or scale according to a Global view, or Local View, etc. Generally, stick with Global view until you get the hang of things.

Layer Selector

Layers are pretty well documented here. In particular, selecting layers to see in is covered in that section on Viewing layers, and moving objects between layers is also discussed in this page. You can also lock the view for this particular window so that it does not change, no matter what you do in other 3D views.

Render Button

The â← Renderâ← button renders an OpenGL copy of the 3D view. Itâ← s pretty much exactly what you see minus the grid, axes, etc., it just shows the objects. It uses the same Draw mode as the 3D view, so itâ← s rather useful if someone asks to see the wireframe of an object youâ← re working on.

CtrlLMB or â § ShiftLMB clicking the button will render an animation of the 3D View, making it useful for making preview renders of animations. The animation will be saved in the rendered Pics folder (Scene context [F10], Output panel, top filespec) in the format of an avi file or image sequence, depending on the format you have chosen (Scene context [F10], Format panel), for the number of frames specified in the Start and End fields (Scene context [F10], Anim panel).

Using the 3D Window

Description

Your 3D window pane is just like a window looking in on a 3D world. To help keep you oriented as to which way is up (Z), an XYZ axis orientation indicator is in the bottom left hand corner, along with the number of the frame youâ ← re working on and the name of the active object. The rest of the view, if in one of the normal orthogonal views (Top, Front, or Side), shows a grid. The interval between each line in the grid is one Blender Unit (BU), an arbitrary unit of measurement. If you are modeling something from the real world, you can set (in your mind) a BU equal to whatever unit of measure your country and culture favor at the moment. If you are a Swiss CERN physicist, then perhaps Angstroms are your thing. If you are a German Engineer, then Millimeter might be in order. If you are an Amsterdam Space Cadet, then Astronomical Units might light up your fancy.

In this 3D space, the active object is highlighted in pink. The active object is always the last object of a number of selected objects and is displayed in a slightly brighter shade that the objects selected with it. The 3D cursor is a red-white circle with a scope crosshairs. LMB clicking moves the 3D cursor to the selected location in 3-dimensional space. Use the Snap key (â § ShiftS) to align the 3D cursor with the grid, the selected object or the active object.

Clicking with the RMB selects the indicated object unless it is already selected, in which case RMB clicking does nothing. â § ShiftRMB button selects another object and keeps the first one(s) still selected, allowing you to select multiple objects. If you â § ShiftRMB -re-click on an already selected object, you will de-select it (like a toggle).

Edit Mode

Enter Edit mode by the mode selector, or by pressing â Tab in the window. In Edit mode, you select the components of the object, and do things to them. Strange, horrible things. Unless youâ ← re good at this stuff and are willing to put in a lot of practice, in which case youâ ← ll get better. Some objects, like cameras, cannot be edited. Press â Tab again to return to the mode from whence you came.

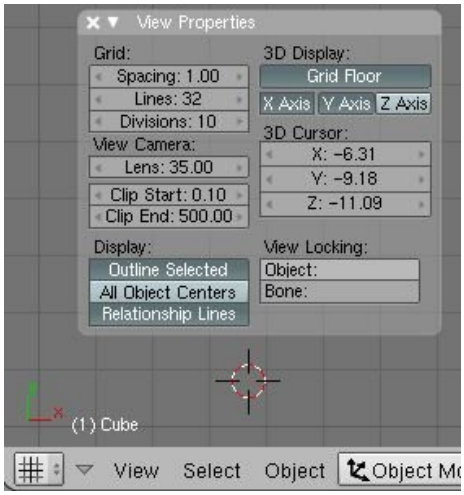
3D Window Toolbox (Popup Menu)

Pressing Space in the 3D window pops up a very handy little menu. You can also press LMB or RMB for one second to do so.

Add

Use this option to add objects to your scene. You can get to this option by:

- Clicking the Add menu item in the User Preferences header at the top of your screen, orâ ←¦
- Pressing Space when your mouse cursor is hovering over any 3D View window.



View Properties panel.
The object, when it is added, is placed wherever your 3D cursor is. It is no more automatically placed in Edit mode (as it used to be), but you can enable this behavior with the Switch to Edit Mode button, within the Edit Methods panel of the User Preferences window. To set the location of your 3D cursor, you can:

- LMB click in a 3D View window, and the cursor jumps to that spot. To place the cursor in 3D space, you will have to click in two windows that have different points of view (once for example in the top view to establish the XY location, and then again in a front view to establish the Z location), orâ ←¦
- Use the View â View Properties window and enter an *exact* XYZ location in the 3D Cursor fields, orâ ←¦
- Select an object whose center is where you want the 3D cursor to be, and â § ShiftSnap the Cursor -> Selection, and the cursor will jump to the selected objectâ ← s location. Note that an object added there will be put inside or will have its surface mushed with the other object (the two objects will attempt to occupy the same space).

Blender supports many different primitives which can be added to your scene. IMPORTANT: If you add an object in Object mode, the primitive is added as a separate object to the scene. If you add an object while in Edit mode for another object, the primitive is added to the other object, forming a compound object from many primitives:

- Mesh - This submenu allows you to add polygonal meshes to your scene. The complex Monkey (called Suzanne) is quite useful for testing purposes (trying out new materials, fur, etc.).
- Curve - This submenu allows you to add curves or paths to your scene. These are useful for modeling curved objects (roller coasters, legs of fancy furniture, etc.) or for making curves for animated objects to follow.
- Surface - This submenu allows you to add NURBS objects to your scene.
- Meta - This submenu allows you to add meta objects to your scene. These are algorithmically-generated objects, and are quite useful for special effects (using two metaballs to animate a cell splitting in half, for example).
- Text - This command allows you to add text objects to your scene.
- Empty - This command allows you to add empties to your scene. Empties arenâ ← t rendered; rather, they are often used to control aspects of other objects. For example, you could use an Empty to control the rotation of an Array modifier. Since they donâ ← t render, you can use them for all sorts of things such as hooks, etc. that require an object to get a location, rotation and/or size from.
- Group - This submenu allows you to add copies of any groups you have in your scene. This is quite useful if you have entire objects grouped together, as you can easily add copies of them.
- Camera - This command adds a camera to your scene.
- Lamp - This submenu contains various types of lamps you can add to your scene. For more information on the different types of lamps, see the lamp documentation in the manual.

- Armature - This command adds an armature (skeleton) to your scene. These are primarily used for animating arms, legs, etc., though if you are in a Pixar-y mood, you can always rig up a lamp!
- Lattice - This command adds a lattice to your scene. These objects do not do anything themselves; however, you can use them to deform objects. In order to have a Lattice deform an object, you have to add a Lattice modifier to the object you want to deform. You can then place the lattice around the object (like a cage), and any changes to the lattice will deform the object. The more detailed the object you are deforming is, the better it will look, so a Subsurf modifier may be helpful here.

Edit

- Enter Editmode - This command will enter Edit mode, which allows you to edit the vertices, edges, and faces of a specific object, rather than manipulating the entire object.
- Duplicate - This command makes a separate duplicate of the selected object(s).
- Duplicate Linked - This command makes a duplicate of the selected object; however, the ObData datablock is shared, so the objects share the same mesh (or curve, â←¦), Ipo curve, material, etc.
- Delete - This command deletes the selected object(s).
- Object Keys - This submenu contains commands related to keyframes.
 - ♦ Show/Hide - This command toggles showing wireframe (ghost) version of the keyframes for the selected object. This is quite useful, as it allows you to visualize the path of the selected object.
 - ♦ Select Next - This command selects the next keyframe for the selected object.
 - ♦ Select Prev - This command selects the previous keyframe for the selected object.

Select

This sub-menu contains exactly the same options as the Select menu in the header.

Transform

- Grab/Move - This command allows you to freely grab (move/translate) the selected object(s).
- Grab/Move on Axis - This submenu contains commands that allow you to move an object along a specific axis: X Global, Y Global, etc.
- Rotate - This command allows you to rotate an object around the view Z axis (i.e. it turns clockwise/counterclockwise around the screen; the rotation axis goes straight into the display).
- Rotate on Axis - This submenu allows you to rotate an object around a specific axis: X Global, Y Global, etc.
- Scale - This command scales the selected object(s).
- Scale on Axis - This submenu contains commands that allow you to scale an object on a certain axis: X Global, etc.
- ObData to Center - This command moves the mesh vertices/curve control-points/etc. of an object so that they are centered according to the current object center.
- Center New - This command moves the center of the object to the center of the object data.
- Center Cursor - This command moves the center of the object to the current location of the 3D cursor.
- Properties - This command toggles the Transform Properties floating panel, which allows you to input exact locations for vertices, as well as location, rotation, and size for entire objects.
- Mirror - This command allows you to mirror (flip) the selection along the appropriate axis. As with all axis-related stuff, â← Globalâ← refers to the scene in general, while â← Localâ← refers to the axis

specific to that object (depending on its rotations).

- Snap - This submenu contains commands that allow you to snap the 3D cursor and the selection to the grid and each other.
 - ♦ Selection -> Grid - This command snaps the selection to the nearest point on the grid.
 - ♦ Selection -> Cursor - This command snaps the selection to the location of the 3D cursor.
 - ♦ Selection -> Center - This command snaps the selected object(s) to the center of the selected object(s). This is most useful in Edit mode, as it allows you to snap vertices, edges, or faces to the center of the object youâ← re working on.
 - ♦ Cursor -> Selection - This command snaps the 3D cursor to the center of the selected object(s).
 - ♦ Cursor -> Grid - This command snaps the 3D cursor to the nearest point on the grid. This is quite useful if you manually clicked to position the 3D cursor, and it didnâ← t land exactly where you wanted.
 - ♦ Cursor -> Active - This command snaps the 3D cursor to the center of the active object.
- Clear/Apply - This submenu contains commands that allow you to clear (reset) or apply (make permanent) the location, rotation, scale, deformation, or duplicates of the selected objects.
 - ♦ Clear Location - This command clears (resets) the location of the selected object(s) to (0,0,0).
 - ♦ Clear Rotation - This command clears (resets) the rotation of the selected object(s).
 - ♦ Clear Scale - This command clears (resets) the scale of the selected object(s).
 - ♦ Apply Scale/Rotation - This command applies the scale and rotation. The object data (mesh vertices/curve points/etc.) is modified so that the scale is 1 and the object isnâ← t rotated at all.
 - ♦ Apply Visual Transform - ???
 - ♦ Apply Deformation - This command applies the deformations to the object. *Note: This is now only available through the Apply buttons in the Modifiers panel (Editing context).*
 - ♦ Make Duplicates Real - This command makes the duplicates (from using DupliVerts or DupliFrames) real objects (so you can edit them individually).

Object

This sub-menu contains a sub-part of the options of the Object menu in the header.

View

This sub-menu contains a sub-part of the options of the View menu in the header.

Render

- Passepartout - This option toggles the Passepartout option of the selected camera. When turned on, in Camera view, it darkens the area around the camera, allowing you to focus on the area thatâ← s actually going to be rendered (is in the cameraâ← s field of view).
- Set Border - This option allows you to drag to select a specific area of the camera view to render. This is useful if youâ← re tweaking a specific detail of an object and donâ← t want to render the entire scene (if youâ← re tweaking an entire object, Local View may be more of what youâ← re looking for). If you want to remove this clipping region from your future rendering, uncheck the Border button (checked by default when you use the Set Border option) in the bottom of the Render panel, Render sub-context, Scene (F10) context.
- Render - This option renders the current scene. Itâ← s the same as the big RENDER button in the Scene context, and can also be activated with F12.

- Anim - This option renders an animation using the current animation settings. Itâ← s the same as the ANIM button in the Scene context, and can be activated by CtrlF12.
- Preview - This toggles the Preview floating panel, which displays a preview (non-antialiased) version of whatever portion of the 3D view is currently underneath it.

The Camera

Add a new camera

Mode: Object mode

Hotkey: â § ShiftA to add new, F9 to change settings.

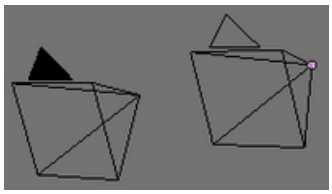
Menu: Add â Camera

In Object mode simply press Space and in the popup menu, choose Add â Camera. New cameras are directed in parallel with the current viewport.

Change active camera

Mode: Object mode

Hotkey: Ctrl0 NumPad



Active camera (left one).
The *active* camera is the camera that is currently used for rendering. Select the camera you would like to make active and press Ctrl0 NumPad (by doing so, you also switch the view to camera view). In order to render, each scene **must** have a camera.

The active camera is the one with the filled â← upâ← triangle on top seen in the 3D viewport. The left camera in the picture.

The active camera, for rendering purposes, is tied to the â← Locks Active Camera and layers to Sceneâ← button (lock icon) for each 3D view port. When a camera is designated as â← activeâ← in a *locked* 3D view, it becomes the rendering camera, and *is made active in all other locked 3D views* (this is a scene-wide setting). When a camera is designated as â← activeâ← in an unlocked 3D view, *it is only â← activeâ← in this viewport* (this is a viewport-wide setting), it has no effect on other 3D views, or on rendering camera. Note that, as the tool-tip of the â← lock-buttonâ← says, this works the same for the active layers selectionâ←¦

Move active camera to view

Mode: Object mode

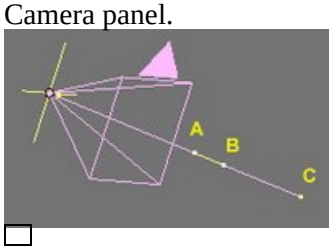
Hotkey: CtrlAlt0 NumPad

Moves the selected camera to current 3D view. Select a camera and then move around in 3D view to a desired position and direction for your camera. Now press CtrlAlt0 NumPad and your selected camera positions itself at your spot, and switches to camera view.

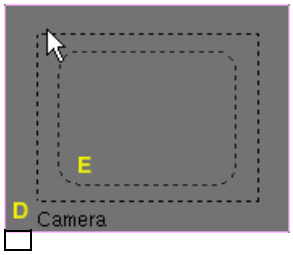
Camera Settings



Camera panel.



Camera.



Camera View.

Cameras are invisible in a scene. They are never rendered, so they donâ← t have any material or texture settings. However, they do have Object and Editing setting panels available which are displayed when a camera is the selected (active!) object.

- Lens, D - Represents the lens, in mm. By enabling the D option, you can specify the â← lensâ← in degrees. When Orthographic is selected, all this will change to a Scale variable.
- DoFDist - Distance to the point of focus. It is shown as a yellow cross on the camera line of sight. Limits must be enabled to see the cross. It is used in combination with the Defocus compositing node.
- Orthographic - Toggles Orthographic mode for rendering. See this page for a more detailed description on Orthographic perspective. You can limit what is rendered by moving the camera. Objects behind the cameraâ← s XY plane are not rendered. When enabled, the Lens size field changes to a Scale field, which includes more/less area in the render.
- Clipping Start/End - Sets the clipping limits. Only objects within the limits are rendered. If Limits is enabled, the clipping will be seen as two yellow dots on the camera line of sight. **C** on the (*Camera*) picture. The first is at cameraâ← s origin.
- Shift X/Y - Shifts the camera viewport (useful to render a low-angle picture of a building, without too much deformations, for example).
- Limits - Toggles viewing of the limits on and off.
- Mist - Toggles viewing of the mist limits on and off. The limits are shown as two white dots on the camera line of sight. **A** and **B** on the (*Camera*) picture. The mist limits are set in the World panel, World sub-context (F8), Shading context.
- Name - Toggle name display on and off. **D** on the (*Camera View*) picture.
- Title Safe - When this is enabled, an extra dotted frame is drawn inside the camera viewport, delimiting the area considered as â← safeâ← for important things, like titlesâ←¦ Shown beside **E** in (*Camera View*).
- Passepartout, Alpha - This mode darkens the area outside of the cameraâ← s field of view, based on the Alpha setting.

- Size - The draw size of the camera in the 3D view. This doesnâ ← t affect the cameraâ ← s output (render), it is just a convenience to enable easier selection of the camera object in the viewport (the camera object can also be scaled using the standard S transform key).

Note
You might be interested to know that a 3D View window, in the Blender user-interface, also has some of these settings, such as Clip Start/End. These of course have nothing to do with cameras, but they serve the same purposes with regard to the 3D window itself: they determine what you can â ← seeâ ← through that particular 3D window, and, generally speaking, how the window behaves (when it *isnâ ← t* in Camera View mode). These settings are accessed through the View â View Properties... menu of the window in question.

Camera Navigation

To control the camera while viewing through it (0 NumPad):

Roll, Pitch, Dolly, and Track

To perform these camera moves, the camera must first be *selected*, so that it becomes the active object (while viewing through it, you can RMB on the solid rectangular edges to select it). Having done so, we can now manipulate the camera using the same commands that are used to manipulate any object:

Roll: Press R to enter object rotation mode. The default will be to rotate the camera in its local Z-axis, which is the definition of a camera â ← rollâ ← .

Pitch: This is just rotation along the local X-axis. Press R to enter object rotation mode, then X X (that is, press the same key *twice*). The first X (or Y or Z for other axis) selects the *global* axis; pressing the same letter a second time selects the *local* axis.

Dolly: To dolly the camera, press G then MMB, LMB to complete.

Track Camera: Press G(rab) and move the mouse (LMB to set position).

Aiming the camera in Flymode

Press â § ShiftF to enter Camera Fly Mode, then move the mouse around to aim the camera. You can also travel it forward/backward using respectively WheelUp/WheelDown. LMB to set (accept) the new orientation, and as usual RMB or Esc to cancel.

Layers

Mode: Object mode

Panel: Object context â Draw

Hotkey: M

Menu: Object â Move to Layer...

Description

3D scenes often become exponentially more confusing with growing complexity. Also, sometimes the artist needs precise control over how individual objects are lit, and do not want lights for one object to affect nearby objects. For this and other reasons below, objects are placed into one or more â← layersâ← . Using Object Layers, you can:

- By selecting certain layers in the 3D View header bar, only objects on those layers are displayed at any one time in your 3D View, speeding up refresh/redraw, reducing virtual-world clutter, and enhancing workflow velocity.
- By making a light illuminate only those objects on its layer, then you control which lights illuminate an object, and vice versa.
- Since particles are affected by forces and effects on the same layer, you can control which forces affect which particle systems.
- Renderlayers cause the rendering of objects on certain layers. By using Renderlayers, you control which objects on currently selected layers are rendered, and which properties/channels are made available for compositing.

Armatures can become very complex, with different types of bones, controllers, solvers, custom shapes, and so onâ←¦ All within a fairly close space, this can also become cluttered. Therefore, Blender provides layers just for armatures. Armature layers are very similar to object layers, in that you can divide up an armature (rig) across layers and only display those layers you wish to work on. Layers (armature and object) behave the same way, armature layers are discussed in the Armatures section.

3D layers differ from the layers you may know from 2D graphics applications: they have no influence on the drawing order and are there (except for the special functions listed above) mainly to allow you to arrange, to organize your scene.

When rendering, Blender only renders those layers selected. If all your lights are on a layer that is *not selected*, you wonâ← t see anything in your render except for objects lit by ambient lighting.

Groups and Parenting are other ways to logically group related sets of objects. Please refer to those applicable sections for more info.

Options

Viewing layers

Blender provides 20 layers; you can choose which are to be displayed with the small unlabeled buttons in the header (*A 3D Viewportâ← s layer buttons*). To select only one layer, click the appropriate button with LMB

; to select more than one, hold â § Shift while clicking.



A 3D Viewportâ← s layer buttons.

To select layers via the keyboard, press 1 to 0 (on the main area of the keyboard) for layers 1 through 10 (the top row of buttons), and Alt1 to Alt0 for layers 11 through 20 (the bottom row). The â § Shift key for multiple selection works for these hotkeys too. By default, the lock button directly to the right of the layer buttons is pressed; this means that changes to the viewed layers affect all 3D Viewports. To select only certain layers in one window, deselect locking first.

Locking the View

To the right of the layer buttons in the 3DView header there's a Lock button. It locks the active camera and layers to this particular view (the layers selected and the orientation/camera used to view the 3D scene). Normally, it is locked, and thus any layer or camera change you make in *another* 3D-View window will be reflected in *this* 3D View window. If you disable it, then this window is not locked to those other windows, and any changes you make in other 3D View windows will not affect this particular view. Now with Ctrl+Numpad0 you can change the camera in another 3D view, and this window will not be affected. In this way, you can make different cameras the active camera in their own 3d-view.

Multiple Layers

An object can exist on multiple layers. For example, a lamp that only lights objects on a shared layer could â← beâ← on layers 1, 2, and 3. An object on layers 2 and 3 would be lit, whereas an object on layers 4 and 5 would not. There are many places where layer-specific effects come into play, especially lights and particles. To place an object on multiple layers, press M and then click â § ShiftLMB the layers you want it to be on (see bellow).

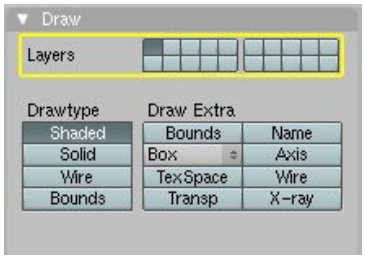
Moving objects between layers



Layer selection.
To move selected objects to a different layer, press M, select the layer you want from the pop-up dialog, then press the OK button. Objects can be on more than one layer at a time. To move an object to multiple layers, hold â § Shift while clicking. If you wish to clone-display the object to additional layers, be sure to â § ShiftLMB click the original layer as well.



Object context selection.
Another way to view or change a selected object layer is via the Draw panel, in the Object context (F7), Object sub-context, as shown opposite.



Layers in Object â Draw panel.
You will then see the layer buttons in the Draw panel, as before the object can be displayed on more than one layer by holding â § Shift while clicking. In this way, you can control what objects are shown on each layer. Layers are useful in keeping your display from being too cluttered, from refresh/redraw time being too long, and to give you control over what to render for later compositing.

Animating Layers

An Objectâ← s layer â← membershipâ← can be animated. E.g. to have objects suddenly appear or disappear in a scene.

Example of Object layer arrangement

As a suggestion, use the top row of layers for the real important things, and the bottom row for those you donâ← t use or change often, or for alternatives for the top row. In a staged set involving mainly two actors, then, you might have, for layers:

1. Lead Actors.
2. Supporting Actors.
3. Supporting Crew (background actors).
4. Particles and effects (vortex, wind).
5. Main Stage.
6. Main backdrops and panels.
7. Main props (tables, chairs).
8. Little props, fillers, decorations, trappings.
9. Cameras, Lights.
10. Lead Actorsâ← armatures.
11. Supporting Actorsâ← armatures.
12. Crew armatures.
13. Alternative clothing.
14. Mesh WIP.
15. Different stage setup, dimensions.
16. Different backdrops that maybe we should use.
17. Other big props that maybe clog up the scene.
18. Props WIP.
19. Additional lighting.

Layer Naming Script



Layer Manager Scripts:

There are also a few scripts available that allows you to give names to layers.

Links:

- 4mm Layer Manager Script.
- Manager Script.

Local or Global View

Mode: Object mode or Edit mode

Hotkey: / NumPad

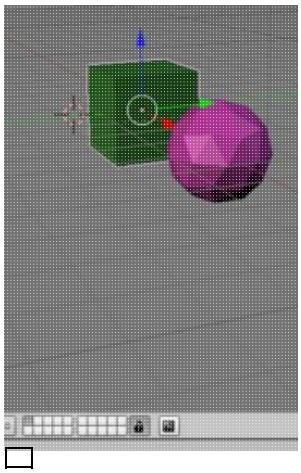
Menu: View â Local View or Global View

Description

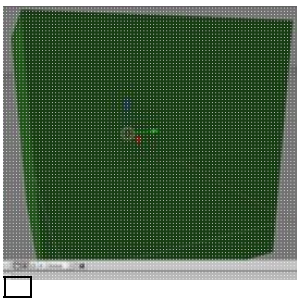
Toggles between Local and Global view mode. The currently selected object is the focus for the mode. An object must be selected to enter Local view mode.

Examples

The layer buttons on the 3D View header disappear while in Local view mode.



Global View.



Local View.

In (*Global View*), the layer buttons are visible and the green cube is selected. (*Local View*) shows that the cube has been focused and centered in the 3D View and the layer buttons have disappeared. This feature is handy when you want to focus on just the object and nothing else. If a scene has thousands of objects visible, it can potentially speed interactivity up because it is the only object visible.

Introduction

This page is based on the grease pencil release logs for blender 2.48.

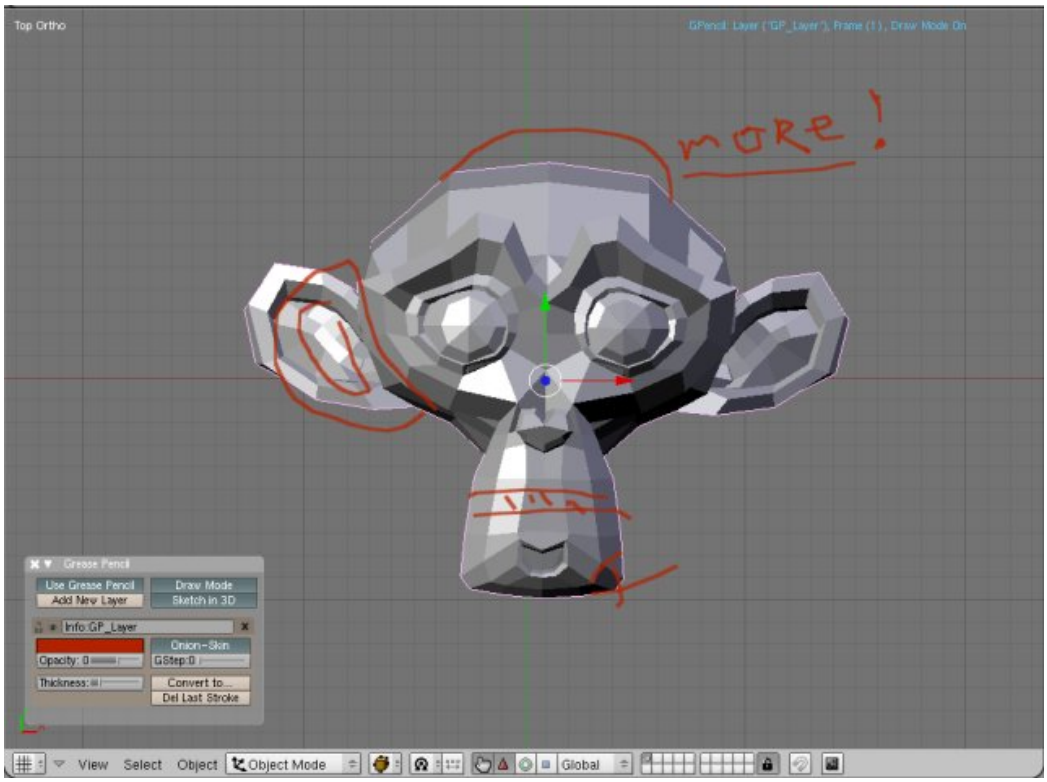
The ability to draw in and/or on viewports using freehand strokes to form sketches and annotations has many benefits for collaborative communication and planning. This can be linked back to traditional 2D-workflows with pencil and paper, where rough â← guidelineâ← sketches were often used for planning and also for communicating ideas quickly. It is often useful to be able to directly scribble on to a work in progress, instead of having to do so in a separate place (i.e. another part of the window, or even in a different application altogether).

What is Grease Pencil?

The name â← Grease Pencilâ← is derived in homage to the wax crayons/pencils that early CG Animators used to draw arcs and other planning notes on their CRTâ← s with.

In addition to uses for animators in planning their poses and motion curves, Grease Pencil can also be useful in a number of scenarios, including but not limited to:

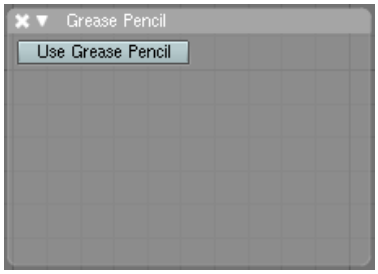
- Planning topology and/or layout of models.
- Directorâ← s shot review tool.
- â← Whiteboardâ← and assignment review tool for educators.



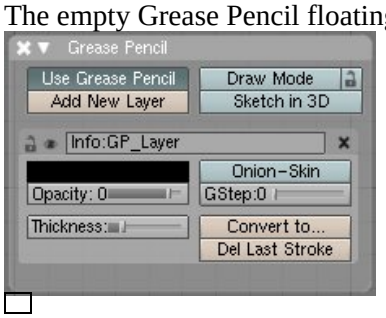
The Grease Pencil in action.

The next few pages explain how to use this tool.

Preparing



The empty Grease Pencil floating panel.



Now all options are visible.

1. The first step when using Grease Pencil is to enable the display of Grease Pencil drawings for the relevant view. To do this, locate the Grease Pencil... entry in the View menu; click on the Use Grease Pencil toggle, in the Grease Pencil floating panel that appears.
2. At this point, click on Add New Layer to add a new layer to draw on. This step is not necessary when you are starting off a new drawing (as a new layer will automatically be created for you), unless you want to customize the line width, colour, and opacity, before drawing. However, if you want to draw on a new layer once layers already exist, it is necessary to click on the button.

Grease Pencil is available for the following window types: 3D View, Node Editor, UV/Image Editor, and Video Sequence Editor.

Warning
Please note that **Grease Pencil data is currently unique to the particular screen area that it was created for**, which means that if **you remove that view, you lose your grease pencil data**.

Drawing

Quick Usage (For Just A Few Strokes)

- **To draw a stroke:** While holding â § ShiftLMB , start dragging the mouse to draw a new stroke. The stroke will finish when you release the mouse button.
- **To erase stroke(s):** While holding AltRMB , start dragging the mouse to erase segments of strokes that fall within the radius of the eraser â← brushâ← .

Note for GNU/Linux users
Under some window managers (like Metacity, the default Gnome one), the Alt key is by default used to move/display the main window â← menuâ← . Which means that you wonâ← t be able to erase strokes with AltRMB dragging. You have two solutions: either to enable Draw Mode (see below), or to modify your Gnome settings (e.g. by using the â← Superâ← /â← Windowsâ← key instead of the Alt one, in general, in something like the System â Preferences â Windows menu).

Easier usage (for drawing more complex sketches)

1. Enable the Draw Mode toggle in top right-hand corner of the Grease Pencil panel.
2. As for quickly drawing a few strokes, use the same mouse buttons to draw and erase, *but* without needing to use the modifier keys too (i.e. LMB to draw, RMB to erase).

Special Tricks in â← Draw Modeâ←

- **Drawing a straight line:** Hold Ctrl while dragging with the LMB to draw a straight line. Although a wavy line will still appear on screen, only the endpoints of that stroke will be used for the final stroke that gets stored. This is a useful feature for architectural uses.
- **Drawing a dot:** Simply click on a spot. This is mentioned here because it is not available when Draw Mode is not enabled.

For Tablet Users

- The thickness of a stroke at a particular point is affected by the pressure used when drawing that part of the stroke.
- The â← eraserâ← end of the stylus can be used to erase strokes too.

Sensitivity When Drawing

The default settings for the sensitivity to mouse/stylus movement when drawing, have been set so that there

shouldnâ← t be too much jitter while still allowing for fine details to be made. However, sometimes these settings may not be appropriate, in which case, the defaults can be found in the User Preferences window, in Edit Methods â← panelâ← .

- Manhatten Distance: The minimum number of pixels the mouse should have moved either horizontally or vertically before the movement is recorded. Decreasing this should work better for curvy lines.
- Euclidean Distance: The minimum distance that mouse should have traveled before movement is recorded. Imagine this as length of imaginary string between last recorded position and mouse-cursor.
- Eraser Radius: This is self-explanatory. It is simply the size of the eraser â← brushâ← , so changing this will affect how likely strokes are going to be covered within the eraser brush and thus erased.
- Smooth Stroke: This turns on the post-processing step of smoothing the stroke to remove jitter. It is only relevant when not drawing straight lines. By default this is off, as it can often cause â← shrinkingâ← of drawings, which is sometimes not that desirable.

Additional notes

- When Draw Mode is enabled, many of the other events that are attached to the LMB and RMB are blocked.
- If you have swapped mouse buttons (in User Preferences window â View & Controls panel â Select with: buttons block), this has no effect on the mapping of mouse-buttons to drawing/erasing operations. However, it may become difficult to select using â § ShiftLMB in such a situation, in which case the tiny â← Lockâ← icon beside the Draw Mode button should be enabled to help alleviate the problems (that will simply disable drawing from occurring with â § ShiftLMB).

Drawing Planes

Sketches are only relevant for the view/view-angle (referred to here as the â← drawing planeâ←) that they were drawn at. There are several different options for how individual strokes (determined by the settings in use when the stroke was created) will be displayed.

- **Screen-Aligned:** This is the default drawing plane for the 3D view, and is also the viewing plane that gets used for the other editors when Stick to View is disabled. All new strokes that get drawn in this mode appear to be â← stuckâ← to the screen-surface (much like markers on a monitor), and will therefore stay unaffected by zooming/translating/rotating the view.
- **View aligned (default for all 2D Views):** New strokes are affected by manipulating the view. This can be turned on/off using Stick to View option.
- **Drawing in 3D-Space (only available in the 3D view):** When Sketch in 3D is enabled, new strokes are drawn in 3D-space, with the position of their points being determined by the position of the 3D-cursor and the view rotation at the time.

Layers

Grease Pencil sketches are organized in layers, much like those you could find in the GIMP or Photoshop. These layers are not related to any of the other layer systems in Blender, and also do not have an upper limit on the maximum number of layers that can be used. Like the layers in the aforementioned applications, these layers can also be renamed, locked, hidden, and deleted.

Their main purpose is to collect together a bunch of sketches that belong together in some meaningful way (i.e. â← blocking notesâ← , â← directorâ← s comments on blockingâ← , or â← guidelinesâ←). For this reason, all the strokes on a layer (not just those made after a particular change) are affected by that layerâ← s colour, opacity, and stroke thickness settings.

By default, most operations occur only on the *active* layer. The active layer can be identified as the one with the different panel colour (in the default set, an light orangy-brown colour). Clicking on a layer, or changing any of its settings will make it the new active layer.

The active layer can also be identified by looking at the status indicator (in the top right-hand corner of every view with Grease Pencil data being shown).

Animation of the Sketches

Grease Pencil can be used to do basic pencil tests (i.e. 2D animation in flipbook style). Sketches are stored on the frame that they were drawn on, as a separate drawing (only on the layer that they exist on). Each drawing is visible until the next drawing for that layer is encountered. The only exception to this is the first drawing for a layer, which will also be visible before the frame it was drawn on.

Therefore, it is simple to make a pencil-test/series of animated sketches:

1. Go to first relevant frame. Draw.
2. Jump to next relevant frame. Draw some more.
3. Keep repeating process, and drawing until satisfied. Voila! Animated sketches.

Onion Skinning

Onion-skinning (also known as ghosting), is a useful tool for animators, as neighboring frame(s) are lightly drawn by Blender. It allows animators to make judgments about movements, by comparing movement from different frames.

Usage Notes:

- Onion-skinning is enabled per layer by clicking on the Onion Skin button.
- The GStep field controls how many frames will be drawn. When GStep is 0, only the drawing on either side of the current frame will be visible. Otherwise, this field specifies the maximum number of frames on either side of the current frame that will result in a neighboring drawing.

Adjusting Timing of Sketches

It is possible to set a Grease-Pencil block to be loaded up in the Action Editor for editing of the timings of the drawings. This is especially useful for animators blocking out shots, where the ability to re-time blocking poses is one of the main purposes of the whole exercise.

1. In an Action Editor window, change the mode selector (found beside the menus) to Grease Pencil (by default, it should be set to Action Editor).
2. At this point, the Action Editor should now display a few â← channelsâ← with some â← keyframesâ← on them. These â← channelsâ← are the layers, and the â← keyframesâ← are the frames at which the layer has a sketch defined. They can be manipulated like any other data in the Action Editor can be.

All the available Grease-Pencil blocks for the current screen layout will be shown. The Area/Grease-Pencil datablocks are drawn as green channels, and are named with relevant info from the views. They are also labeled with the Area index (which is currently not shown anywhere else though).

Copying Sketches

It is possible to copy sketches from a layer/layers to other layers in the Action Editor, using the â← Copyâ← /â← Pasteâ← buttons in the header. This works in a similar way as the copy/paste tools for keyframes in the Action Editor.

Sketches can also be copied from one screen (or view) to another using these tools. It is important to keep in mind that keyframes will only be pasted into selected layers, so layers will need to be created for the destination areas too.

Converting Sketches to Other Forms

In the 3D view, sketches on the active layer can be converted to geometry, based on the current view settings. Sketches are converted into geometry by transforming the points recorded when drawing (which make up the strokes) into 3D-space (based on the current view settings). Currently, all points will be used, so it may be necessary to simplify or subdivide parts of the created geometry for standard use.

Sketches can currently be converted into one of three types:

- **Armature**: Each stroke is converted into a bone chain, which is assigned to an armature named after the active layer. The bones in each chain are connected and parented to each other. Also, bones inherit their envelope radii from the thickness of their stroke at each recorded point.
- **Bezier Curve** and **Path**: Each stroke is converted into a separate curve within a curve object thatâ← s named after the active layer. Handles are automatically set to be â← freeâ← handles (i.e. the black type), and are set to be in the same places as the control-points. The weight/radius of the curve at each control-point is set to equal the thickness of the stroke at each recorded point. However, in order to see that, you need to open a Transform Properties panel in Edit mode, Weight field.

Introduction

There are many things you can do with the selected object within your virtual world. You can twist it, make it bigger or smaller, spin it around, change its shape, and so on. This section tells you how to do these things to your objects.

As you will find with most of Blender, there isnâ← t just â← one wayâ← to do things. For infrequent users, there is always the context-sensitive menus. For more experienced users, there are hotkeys, where merely tapping a key performs an action.

Menu Options

With an object selected in a 3D view, the menu bar shows the View, Select and Object options. Click the Object one to manipulate the object. The menu that pops up has the option Transform. Hovering over Transform pops up a sub-menu, showing you selections (and hot keys on the right) for manipulating the object:

- Grab/Move
- Rotate
- Scale
- Convert To Sphere
- Shear
- Warp
- Push/Pull

Since there are also hotkeys to do these things, please read on to the next section.

Grab/Move

Mode: Object mode and Edit mode

Hotkey: G

Menu: Object/Mesh/etc. â Transform â Grab/Move

Description

One of the fastest ways to move things in 3D space is with G. Pressing this hotkey will enter the â← grab/moveâ← transformation mode where the selected object or data is moved according to the mouse pointerâ← s location. The distance from the mouse pointer to the manipulated object has no effect.

Options

LMB 
Confirm the move, and leave the object or data at its current location on the screen.

MMB 
Constrain the move to the X, Y or Z axis.

RMB  or Esc
Cancel the move, and return the object or data to its original location.

Wheel 

Object mode

Zoom in/out

Edit mode

- Proportional Editing disabled: Wheel  does zoom in/out
- Proportional editing enabled: only vertices within the circle of influence are moved. In this case scroll the Wheel  to change the radius of the the circle of influence.

See the rest of the Manipulation in 3D space section for additional options available within the transformation modes.

Rotate

Mode: Object mode and Edit mode

Hotkey: R

Menu: Object/Mesh/etc. â Transform â Rotate

Description

Pressing R will enter the â← rotateâ← transformation mode where the selected object or data is rotated according to the mouse pointerâ← s location. This mode uses the angle from the pivot point to the mouse pointer as the angle for rotation, so moving the mouse further from the object/data will give more fine-grained precision (i.e. the mouse movement will affect the rotation less, for the same mouse distance moved).

Options

LMB
: Confirm the rotation, and leave the object or data at its current rotation on the screen.

MMB
: Constrain the rotation about the X, Y or Z axis.

RMB or Esc
: Cancel the rotation, and return the object or data to its original rotation.

Wheel

: Object mode

: : Zoom in/out

: Edit mode

: : · Proportional Editing disabled: Wheel does zoom in/out
: : · Proportional editing enabled: only vertices within the circle of influence are moved. In this case scroll the Wheel to change the radius of the the circle of influence.

R (Trackball)
: Pressing R while already rotating toggles the rotation mode between a single axis rotation (either aligned to the screen or around a certain axis) and a two axis, â← trackballâ← rotation. In â← Trackballâ← mode, the rotation of the object is controlled by both the X and Y location of the mouse pointer, similar to the trackball view rotation mode. This can be a quick way to rotate an object into place, without having to keep changing the view rotation while adjusting. However, it is more difficult to use than the simple rotation.

See the rest of the Manipulation in 3D space section for additional options available within the transformation modes.

Scale

Mode: Object mode and Edit mode

Hotkey: S

Menu: Object/Mesh/etc. â Transform â Scale

Description

Pressing S will enter the â← scaleâ← transformation mode where the selected object or data is scaled inward or outward according to the mouse pointerâ← s location. The object/dataâ← s scale will increase as the mouse pointer is moved away from the pivot point, and decrease as the pointer is moved towards it. If the mouse pointer crosses from the original side of the pivot point to the opposite side, the scale will continue in the negative direction, making the object/data appear flipped. The precision of the scaling is determined by the distance from the mouse pointer to the object/data when the scaling begins.

Options

LMB
: Confirm the scale, and leave the object or data at its current scale on the screen.

MMB
: Constrain the scaling to the X, Y or Z axis.

RMB or Esc
: Cancel the scale, and return the object or data to its original scale.

AltS (only in Edit mode)
: Scales along the selected normal direction.

Wheel

: Object mode

: : Zoom in/out

: Edit mode

- Proportional Editing disabled: Wheel does zoom in/out
- Proportional editing enabled: only vertices within the circle of influence are moved. In this case scroll the Wheel to change the radius of the the circle of influence.

See the rest of the Manipulation in 3D space section for additional options available within the transformation modes.

Precision

Mode: Object mode and Edit mode

Hotkey: Ctrl, â § Shift

Description

Holding Ctrl and/or â § Shift during a manipulation can be used to control the precision more finely.

Options

Ctrl (Snap)
: Snap the transformation:
 - ◊ For grab/move and scale, on 1 blender unit (or 10BU, 1/10 BU, etc., depending of the zooming view).
 - ◊ For rotation, on 5 degrees step.

â § Shift (Precise)
: Allow finer transformation control but not by precise values (i.e. by fixed steps).

â § ShiftCtrl
: Snap the transformation on finer steps than with only Ctrl:
 - ◊ For grab/move and scale, on 1/10 of the precision with only Ctrl (e.g. if with Ctrl you have steps of 1BU, with â § ShiftCtrl you have steps of 1/10 BU).
 - ◊ For rotation, on 1 degrees step.

Numeric transformations

Mode: Object mode and Edit mode

Hotkey: G/R/S, then + NumPad/- NumPad and 0 NumPad-9 NumPad or Keyboard 0-9

Description

When doing a transformation, instead of using the mouse (imprecise work), you can directly pass a precise numeric value.

Hit - NumPad (on a French keyboard, the minus is under the number 6 key, so you must use the Numpad) if you want negative values, and then a numeric value. You can see the values in the header of the 3D View. If you were using an axis constraint (global or local), the value is applied to that axis. If there is no axis constraint and you want to choose which axis to control, you can use â Tab to switch from any of the axis (a cursor appears after the value). You can also return to a mouse control with â Backspace.

Clear transformations

Mode: Object mode

Hotkey: AltG, AltS, AltR

Menu: Object â Clear/Apply â Clear Location, Clear Scale, Clear Rotation

Description

You can clear any transformation done in Object mode.

Options

AltG
: Clear the location of the selection.

AltS
: Clear the scale of the selection.

AltR
: Clear the rotation of the selection.

Align to

Mode: Object mode

Hotkey: CtrlAltA

Menu: Object â Transform â Align to Transform Orientation

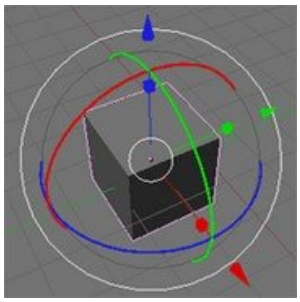
Description

Align selected objects to a specific Transform Orientation.

Manipulators

Mode: Object mode and Edit mode

Hotkey: CtrlSpace



Combination manipulator.
When using the normal Transform commands (G for Grab, R for Rotation, S for Scale), they will work only parallel to the current view. The transform commands can be augmented with additional modifier keys, e.g. pressing X, Y or Z immediately after G, R, or S will constrain the transform to that global axis. Using MMB will do the same but along the nearest axis from the mouse pointer. Pressing R a second time, before pressing LMB (i.e. while the rotation command is in progress), releases the default â← parallelâ← rotation constraint, allowing free rotation in any axis.

Read more about Axis Locking.

Manipulators provide a visual representation of the transform commands and allow movement, rotation and scaling along any axis in any mode of the 3D window.



Manipulator Header.
The manipulator can also be accessed in the header of the 3D View window (â § ShiftLMB can select more than one manipulator at the same time):

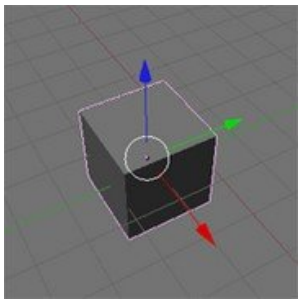
- Hand: Enable/disable the manipulators.
- Triangle: Translation/location.
- Circle: Rotation.
- Box: Scale.
- Transform Orientation: Control the orientation of the axis for transformations.

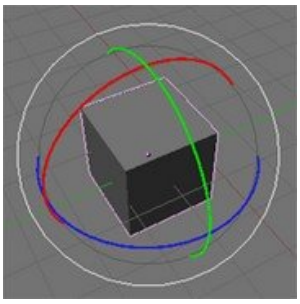
Manipulator types

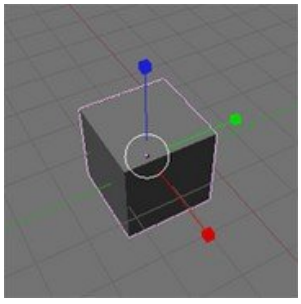
Mode: Object mode and Edit mode

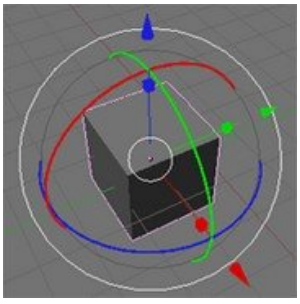
Hotkey: CtrlAltG for Grab/Move, CtlrAltR for Rotate, CtrlAltS for Scale

There is a separate manipulator for each Transform Command. Each can be viewed separately or in a combination.


Translate.


Rotate.


Scale.


Combination.

Other Manipulator controls

Holding down Ctrl constrains the action in certain increments (see this page for details).

Holding down â § Shift when you LMB click on one of the handles, the manipulator action will be performed on the other two axes to the one you clicked on (you can let go of â § Shift once you have clicked).

LMB on the white circle (largest circle around rotation manipulator) do the same as pressing R.

LMB on the grey circle (inner circle around rotation manipulator) do the same as pressing R twice (â← trackballâ← rotation).

Manipulator Preferences



Manipulator preferences.
The settings of the manipulator (e.g. its size) can be found in the View & Controls section of the User Preferences window.

- Size: Diameter of Widget, in 10 pixel units.
- Handle: Size of widget Handles, as a percentage of widget radius (Size/2).
- Hot spot: Hotspot size (in pixels) for clicking Widget Handles.

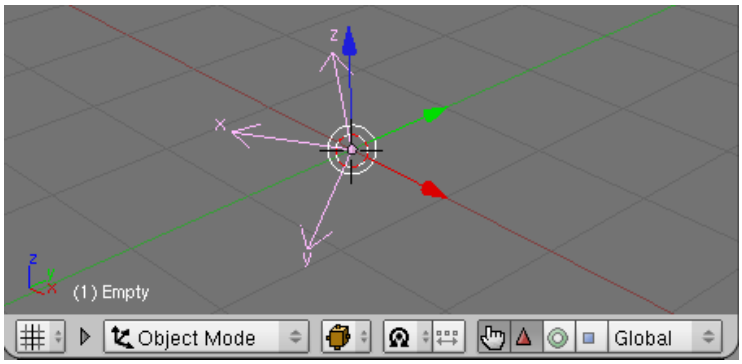
Transform Orientation

Mode: Object mode and Edit mode

Hotkey: AltSpace

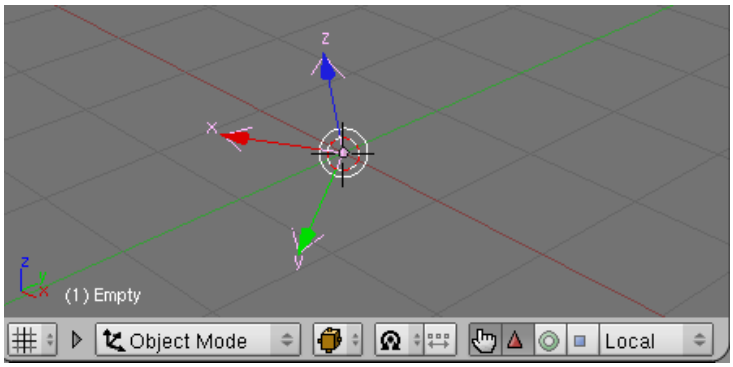
The default (global) transform manipulator is very useful but in some situations itâ← s not (for example: scale a rotated object along the rotated axis). Luckily Blender can change the orientation of the Transform Manipulator.

Below is a list of different transform orientation types. On every image, compare the position of the manipulator axes (color axes over the object) with the global (lower left corner of the 3D window) and local (the object is an empty, so just the local axes of the object are shown) ones.



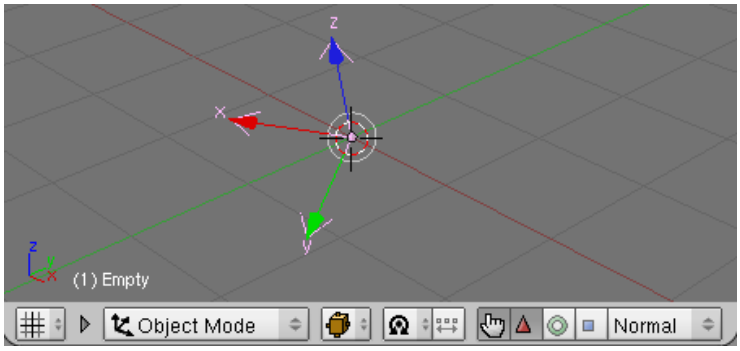
Global.

Global
: The manipulator matches the global axis.



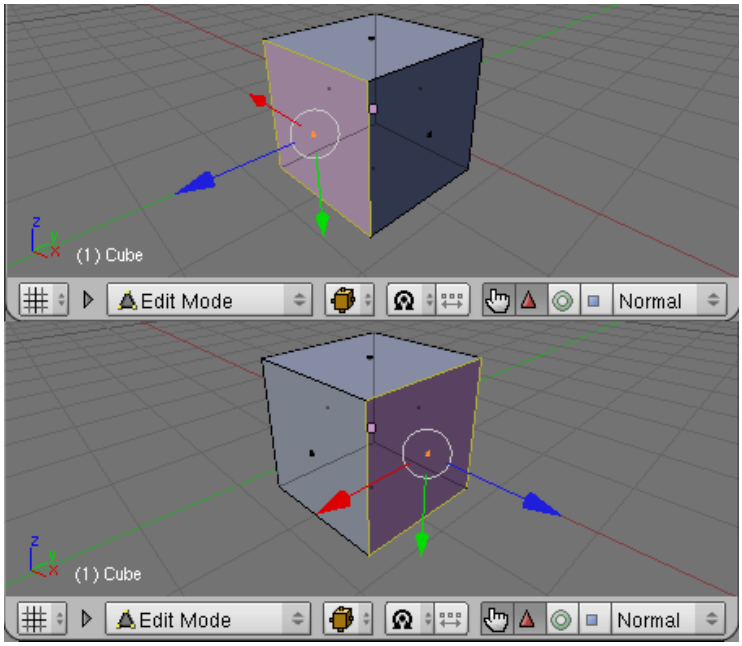
Local.

Local
: The manipulator matches the object axis.



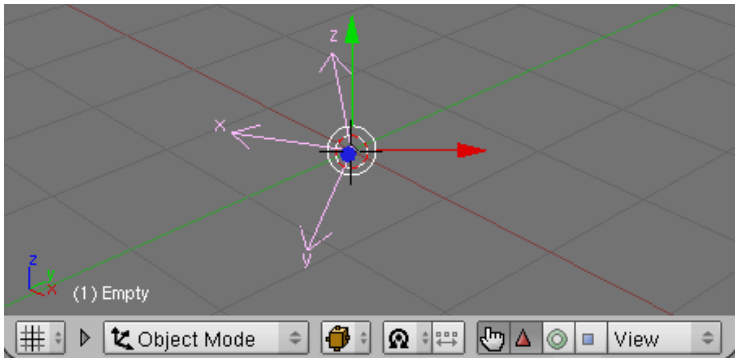
Normal.

Normal
: The Z axis of the manipulator will match the normal vector of the selected object. Not very useful for the empty, see the example below.



Cube example.

A better example using an object with face normals, selecting faces in Edit mode:



View.

View
: The manipulator will match the 3D view, Y â Up/Down, X â Left/Right, Z â Towards/Away from you.

For more precisions about the transform orientations, see this page.

Gestures

Mode: Object mode and Edit mode

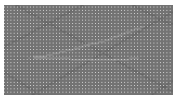
Hotkey: LMB (drag)

Blenderâ← s 3D transform modes can also be invoked by drawing mouse gestures. The tool is designed to figure out what mode to enter based on a hand-drawn gesture. After a gesture is drawn as described below, release the LMB. Move the mouse without pressing any button, then click the LMB when you achieve the effect you want. To cancel, click the RMB, even if movement in the scene has occurred.

There are three gestures the tool recognizes:

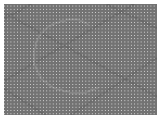
- Scale
- Rotate
- Translate

Scale



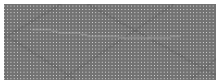
Example â← Scaleâ← gesture.
To activate Scale mode, draw what appears to be a â← **V**â← shaped path using the LMB. (*Example â← Scaleâ← gesture*) is an example of a hand-drawn â← **V**â← . It doesnâ← t have to be exact but the closer and sharper it is the better the tool will understand. If the â← **V**â← has some roundness in it, it most likely will be interpreted as a request for Rotate mode.

Rotate



Example â← Rotateâ← gesture.
To activate Rotate mode, draw what appears to be a â← **C**â← shaped curve using the LMB. (*Example â← Rotateâ← gesture*) is an example of a hand-drawn â← **C**â← . It doesnâ← t have to be exact but the smoother the curve the better the tool will understand. If the â← **C**â← has a sharp corner in it, it most likely will be interpreted as a request for Scale mode.

Translate



Example â← Translateâ← gesture.

To activate Translate mode, draw what appears to be a â← **â**← â← or line using the LMB. (*Example â← Translateâ← gesture*) is an example of a hand-drawn â← **â**← â← . It doesnâ← t have to be exact but the straighter the line the better the tool will understand. If the â← **â**← â← deviates too much from a straight line, it most likely will be interpreted as a request for either Scale or Rotate modes.

Normal axis locking

Mode: Object mode and Edit mode (translate, rotate, scale, extrude)

Hotkey: X, Y, Z

Blender has a very useful option: if you want toâ←¦

- S - Scale,
- R - Rotate,
- G - Move/Grab/Translate,
- E - Extrude,

â←¦the selected object, the operation can be constrained to one axis:

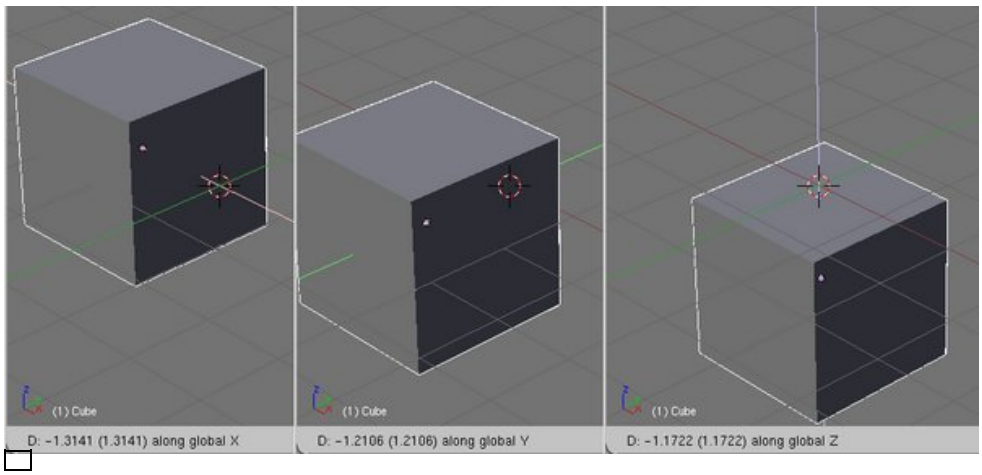
Press the appropriate key to start the operation (e.g. S to scale), then press X, Y, or Z to constrain to the corresponding global axis.

If the same â← axisâ← key is pressed again, the operation is constrained to the objectâ← s local axis (by default : it depends in fact on the selected transform orientation). E.g. G, X, X to constrain to the objectâ← s local X axis.

Move/Translate

Press:

- G, X for the global X-axis (along left-right).
- G, Y for the global Y-axis (along forward-backward).
- G, Z for the global Z-axis (along up-down).

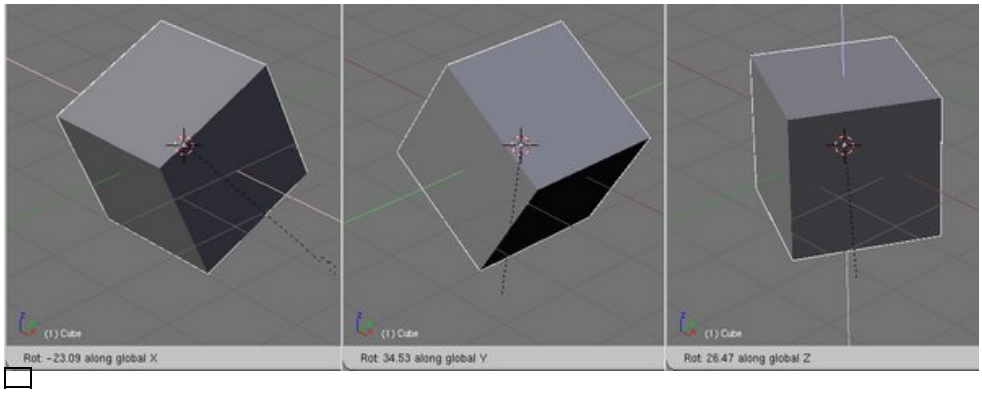


Translation/Grab with locked axis X, Y and Z.

Rotate

Press:

- R, X for the global X-axis (around left-right).
- R, Y for the global Y-axis (around forward-backward).
- R, Z for the global Z-axis (around up-down).

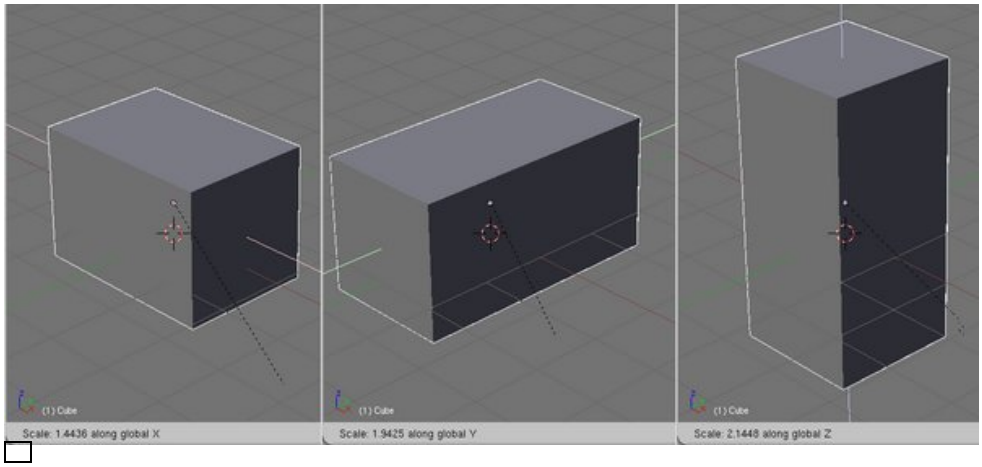


Rotation with locked axis X, Y and Z.

Scale

Press:

- S, X for the global X-axis (along left-right).
- S, Y for the global Y-axis (along forward-backward).
- S, Z for the global Z-axis (along up-down).



Scale with locked axis X, Y and Z.

Plane axis locking

Mode: Object mode and Edit mode (translate, rotate, scale, extrude)

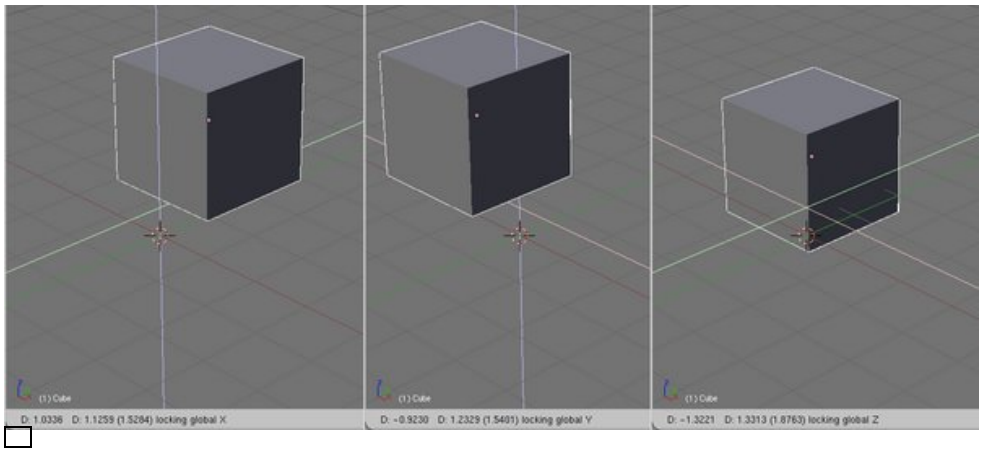
Hotkey: â § ShiftX, â § ShiftY, â § ShiftZ

You can lock two axes at once at both the global and the local axes. Press â § ShiftX, â § ShiftY or â § ShiftZ to lock the *other* two axes:

- â § ShiftX - Uses Y+Z
- â § ShiftY - Uses X+Z
- â § ShiftZ - Uses X+Y

Move/Translate

- Press G, â § ShiftX to move only along the global Y-axis and Z-axis (not left-right).
- Press G, â § ShiftY to move only along the global X-axis and Z-axis (not forward-backward).
- Press G, â § ShiftZ to move only along the global X-axis and Y-axis (not up-down).



Translation/Grab with locked axes Y+Z, X+Z and X+Y.

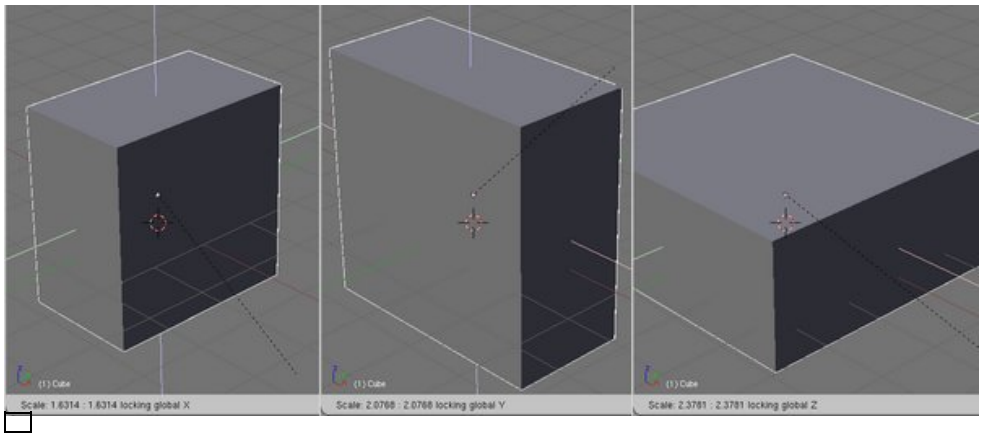
Rotate

For rotation pressing â § Shift has no effect (except for the display of different axes).

Rotating an object by locking two axes has the same effect as using only one axis to lock it, as rotation using two axes will rotate the object on the unlocked axis.

Scale

- Press S, â § ShiftX for the global Y-axis and Z-axis (not left-right).
- Press S, â § ShiftY for the global X-axis and Z-axis (not forward-backward).
- Press S, â § ShiftZ for the global X-axis and Y-axis (not up-down).

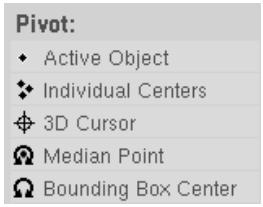


Scale with locked axes Y+Z, X+Z and X+Y.

Pivot

Mode: Object mode and Edit mode

Menu: Droplist in the header of the 3D view



Pivot Point modes.
The pivot point is the point in space around which all rotations, all scalings and all mirror transformations are centered. You can chose among five general modes for your pivot points which can be selected from a drop-down list in the header of any 3D area, as seen here in (*Pivot Point modes*). Your job is to choose the most efficient type for the task and to position pivot point accurately.

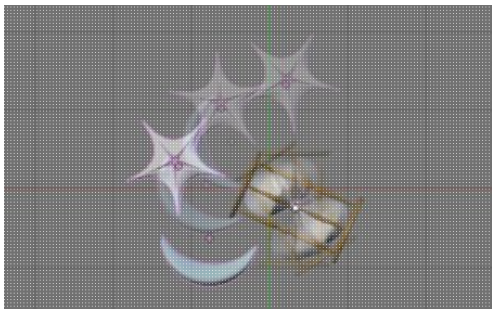
Active Object

Mode: Object mode and Edit mode

Hotkey: Alt.

Forget the name: it is not limited to objects. In Edit mode the *active* element can also be a vertex, an edge or a face.

In Object mode



Rotation around the active Object.
What happens in Object mode is pretty simple: rotations and scalings happen around the active Objectâ← s center. This is illustrated by (*Rotation around the active Object*), where the center of the active Object, an hourglass, is the only thing to remain still.

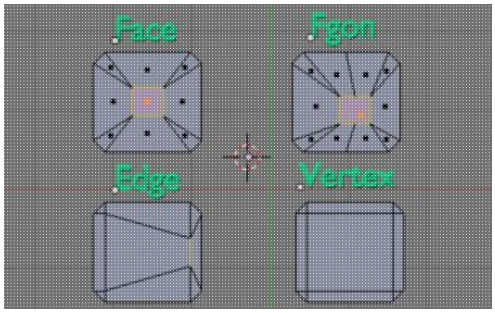
In Edit mode

It may seem like there is a lot of complexity to using the active element as a pivot point in Edit mode but, in reality, the many possibilities result of only a few rules:

- The pivot point is always at the median of the active element(s).
- The transformations occur by transformation of the **vertices** of the selected element(s). If an unselected element shares one or more vertices with a selected element then the unselected one will get some degree of transformation also.

Letâ← s examine the following examples: in each case we will see that the two rules apply.

Single selection

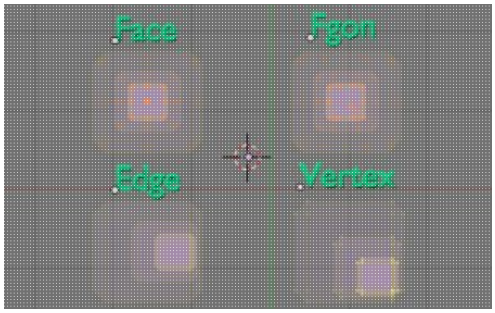


Edit mode and only one element selected.
When one single element is selected it becomes automatically active. In (*Edit mode and only one element selected*), you can see that when it is transformed its vertices move, with the consequence that any adjacent element which shares one or more vertices with the active element is also transformed somewhat.

Lets review each cases:

- *Faces* have their pivot point where their selection dot appears, which is where the median of their vertices is.
- *Fgons* behave the same, but notice that the selection dot can be off compared to faces.
- *Edges* have their pivot point on their middle since this is always where the median of an edge is.
- A single *Vertex* has no dimensions at all so it canâ← t show any transformation (except translation, which is not affected by the pivot pointâ←¦).

Multiple selection



Edit mode and multiple selections.
When multiple elements are selected they all transform.

The pivot points stay in the same place as what weâ← ve seen above, with only one exception for Fgons. In (*Edit mode and multiple selections*), the selected elements have been scaled down from large to small:

- For *Faces* the transformation occurs around the selection dot of the active face.
- *Fgons* behave exactly like faces.
- *Edges* also keep the same behavior with their pivot point at their median.
- There is a case for *Vertices* this time: the active Vertex is where the pivot point resides. All other vertices are transformed relative to it.

Again, as we have seen, all there is to remember is the two rules:

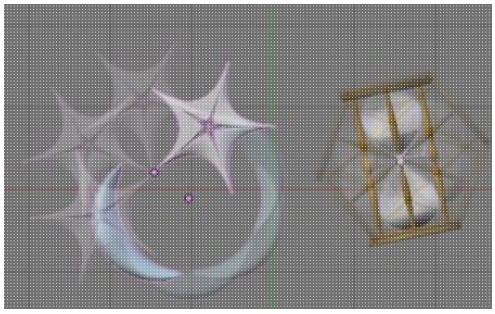
- Pivot point at the median of the active elementâ ← s vertex or vertices.
- All selected vertices, directly or as part of a bigger element, e.g. a face, and only them are transformed.

Individual Center

Mode: Object mode and Edit mode

Hotkey: Ctrl.

In Object mode



□
Rotation around the individual centers.
In (*Rotation around the individual centers*), the object centers of each object remains at the same location while each object is rotating around them.

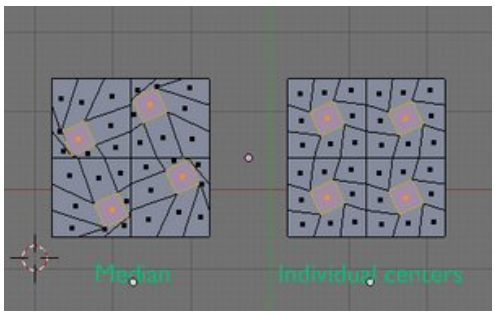
Positioning the center of the objects is a most useful technique that affords you more control over your animations. Letâ ← s examine (*Rotation around the individual centers*):

- The center of the hourglass is coincident with the median of all its components.
- The positioning of the center of the star at the tip of one of its branch allows for a rotation around it with only the use of Rot Ipos that are more faithfully interpolated than what you get using the 3D cursor in the same position.
- The center of the crescent is completely outside of what to us appears to be the object. You must understand that the center really marks **where the object is**; what we see on the screen is a description of what the object is made of: vertices, colors, stuff, and it can very well happen to be off-center, like for the crescent here.

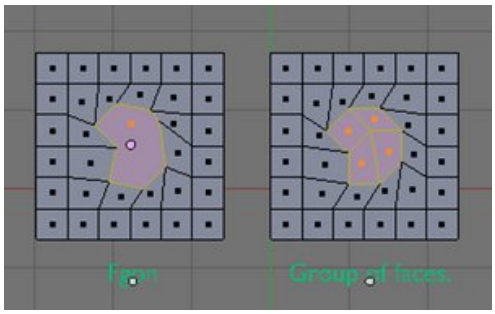
In Edit mode

With the vertex or the edge selection methods in use, a selection of vertices or of edges has its pivot point at the median of the set of vertices so selected. For more information see the Median Point pivot section.

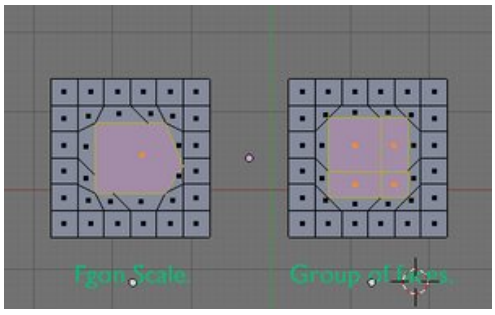
As soon as the face selection method is in use though, the pivot point as the center of those faces becomes possible.



Individual rotation of multiple faces.
It is possible to rotate individually each face only in face selection mode. Only faces that donâ← t touch each other can be transformed in this way without deforming. You can only use the Proportional Editing Tool (PET) while translating individual faces this way.

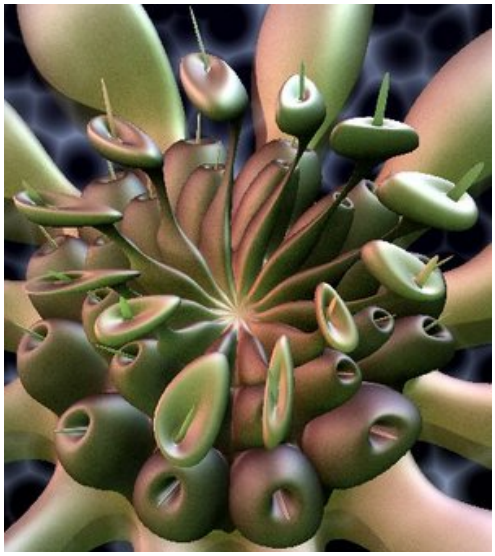


Fgon rotation with individual centers pivot point.
Faces that touch, even when they are inside an Fgon, are deformed when rotated with individual centers as the pivot point.



Problems with Fgon and groups of faces scaling.
Fgons and groups of faces can be scaled and their outside perimeter wonâ← t be deformed. The individual faces inside them arenâ← t uniformly scaled though, something you should take into account.

All those deformations wonâ← t happen if you are not using the face selection mode; it becomes impossible to edit more than one face or one group of faces at a time though.



Modelisation with faces and individual centers pivot point.
Once you are aware of its limitations and pitfalls, this peculiar tool can save a lot of time and lead to unique shapes. This â← anemoneâ← was modeled from a 12 sided cylinder in about 10 minutes by using repeatedly this workflow: extrusions of individual faces, scaling with *median as a pivot point*, and scaling and rotations of those faces with *individual centers as pivot points*.

3D Cursor

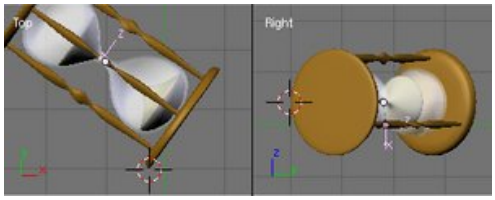
Mode: Object mode and Edit mode

Hotkey: .

The 3D cursor is the simplest most intuitive of all pivot points, it allows for total control of the results. It can be summarize by this: position it and transform.

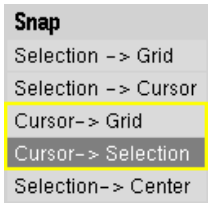
Positioning the 3D cursor

There are a few methods to position the 3D cursor.



Positioning the 3D cursor with two orthogonal views.

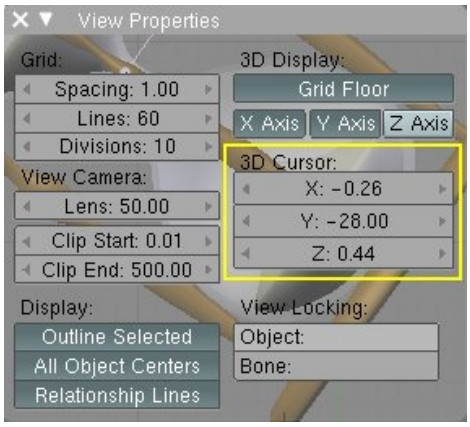
- Using LMB in the 3D area. For accuracy you should use two perpendicular orthogonal 3D views, i.e. any combination of top (7 NumPad), front (1 NumPad) and side (3 NumPad). That way, in one view you can control the positioning along two axes and determine depth in the second view.



The Snap menu.

- Using snaps:
 - â § ShiftS â Cursor -> Grid to send the 3D cursor to the nearest **visible** point of the grid.
 - â § ShiftS â Cursor -> Active to send the 3D cursor to the *active* (last selected) object, edge or face center, or to the *active* vertex.
 - â § ShiftS â Cursor -> Selection to send the 3D cursor to:
 - ◊ The object center of an object.
 - ◊ A vertex.
 - When there is more than one element in the selection and the Bounding Box Center pivot point is selected, â § ShiftS sends the 3D cursor to:
 - ◊ In Object mode, to the **center of the bounding box** surrounding the objects centers.
 - ◊ In Edit mode, to the **center of the bounding box** surrounding the selected vertices (even in edge or face selection mode, it really is the vertices that are selected indirectly that are taken into account).
 - When the Median Point pivot point is selected, â § ShiftS sends the 3D cursor to:
 - ◊ In Object mode, to the median of the object centers.
 - ◊ In Edit mode, to the median of the selected vertices.

Lots of possibilities *is* lots of power!



The View Properties panel.

- Numerically, you can use the View Properties... entry of the View menu (from the 3D view header), and then enter the 3D Cursor location in the 3D Cursor section of the resultant View Properties panel that should now be visible.

Transformation

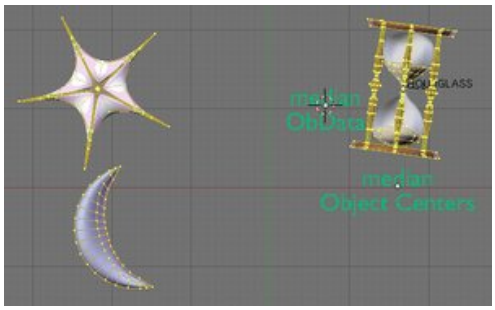
All thereâ ← s left to do is to select the 3D Cursor as the pivot point and rotate, scale or mirror.

Median Point

Mode: Object mode and Edit mode

Hotkey: Ctrl,

We can assimilate the â← median pointâ← to the notion of Center Of Gravity (COG): supposing that every element of the selection has the same mass, the median point would sit at the COG, the point of equilibrium for the selection. This is very abusive as we will see soon enough. Yet it helps predicting where the median point should be when planing a scene.



Median point of object centers and ObData.

In Object mode

For objects, only the object center is taken into account. Moreover, each object center is assumed to have the same mass. This can lead to very counter-intuitive results. On the (*Median point of object centers and ObData*), we see that the median of the objects sits far from the middle of the object centers. That is because the ObData (the geometry) of the moon and the star is way off their object center.

In Edit mode

Still on (*Median point of object centers and ObData*), we see that even the position of the median point for the ObData is surprisingly close to the hourglass: this is because it has much more vertices (611) than the moon (81) and the star (130). Blender supposes that every vertex has the same weight.

Transformation

Once the Median Point has been chosen from the list, the transform manipulator (if enabled) immediately cling to it, giving and excellent visual clue: all the rotations, scalings and mirror will happen around this point.

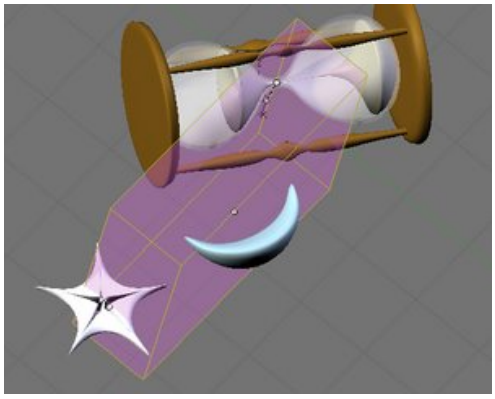
Bounding Box Center

Mode: Object mode and Edit mode

Hotkey: ,

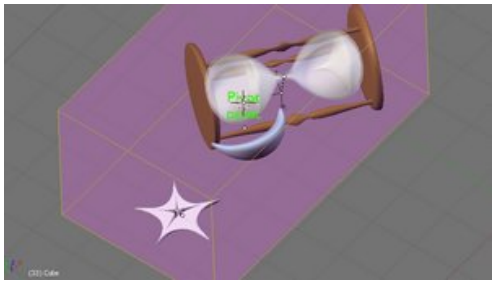
The bounding box is a rectangular box that is wrapped as tightly as possible around the selection. It is oriented parallel to the world axes. In this mode the pivot point lies at the center of the bounding box.

In Object mode



The bounding box in Object mode.
In Object mode, the bounding box is wrapped around the object centers and does not take into account the ObData (geometry).

In Edit mode



The bounding box in Edit mode.
This time it is the ObData that is enclosed in the box because all of its vertices were selected. The bounding box in Edit mode takes no account of the objects centers but only of the selected vertices.

Proportional Edit

Mode: Edit mode

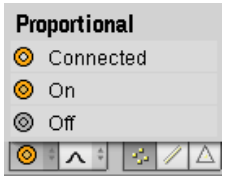
Hotkey: O / AltO / â § ShiftO

Menu: Mesh â Proportional Editing

When working with dense geometry, it can become difficult to make subtle adjustments to the vertices without causing nasty lumps and creases in the modelâ← s surface. When you face situations like this, use the proportional editing tool. It acts like a magnet to smoothly deform the surface of the model, without creating lumps and creases, by also modifying unselected vertices within a given range, not just the selected ones.

Sculpting
Blender offers a complete Sculpt mode which contains a set of brushes and tools for proportionally editing a mesh without seeing the individual vertices.

Options



Proportional Editing icon.



Falloff menu.
The Proportional Editing mode menu is available in Edit mode, on the 3D View header.

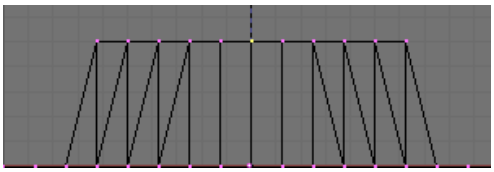
Off (O)
: Proportional Editing is Off, only selected vertices will be affected.

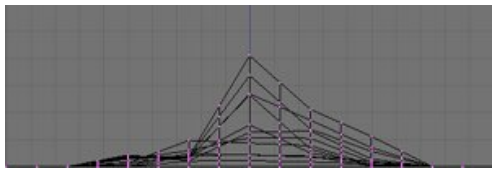
On (O or AltO)
: Vertices other than the selected vertex are affected, within a defined radius.

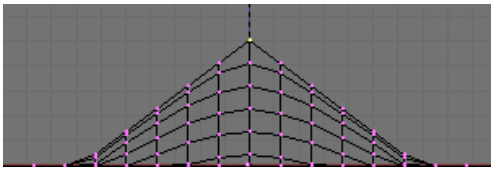
Connected (AltO)
: Rather than using a radius only, the proportional falloff propagates through connected geometry. This means that you can easily proportionally edit the vertices in a finger of a hand, without affecting the other fingers, since although the other vertices are nearby spatially, they are far away following the topological edge connections of the mesh. The icon will be cleared (grey) in the center when Connected is active.

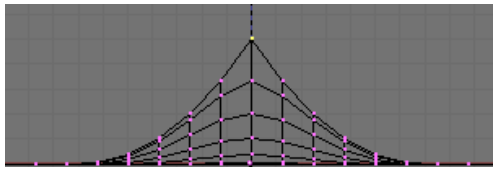
Falloff

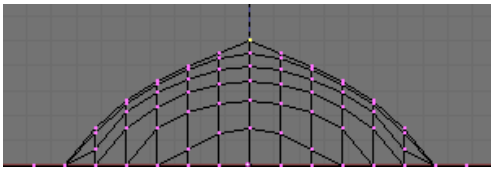
While you are editing, you can change the curve profile used by either using the Mesh â Proportional Falloff submenu, using the toolbar icon (*Falloff menu*), or by pressing â § ShiftO to toggle between the various options.

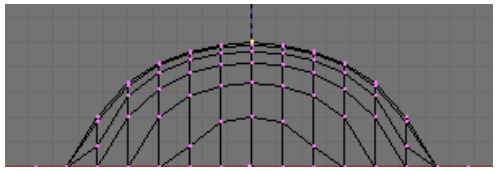

Constant, No Falloff.

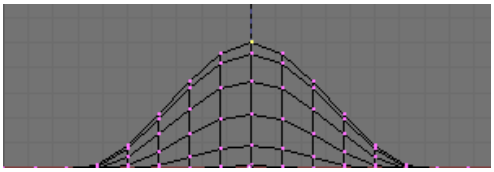

Random Falloff.

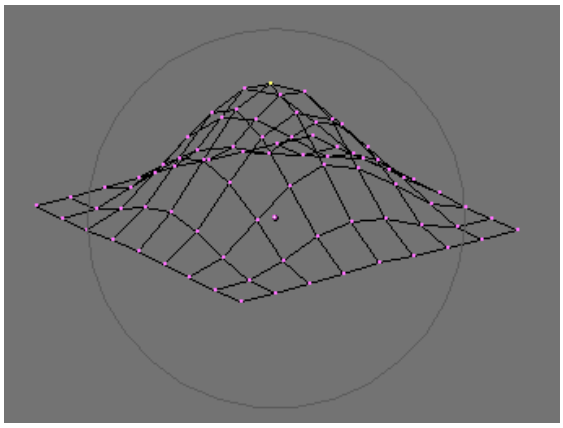

Linear Falloff.


Sharp Falloff.


Root Falloff.


Sphere Falloff.


Smooth Falloff.


Influence circle.

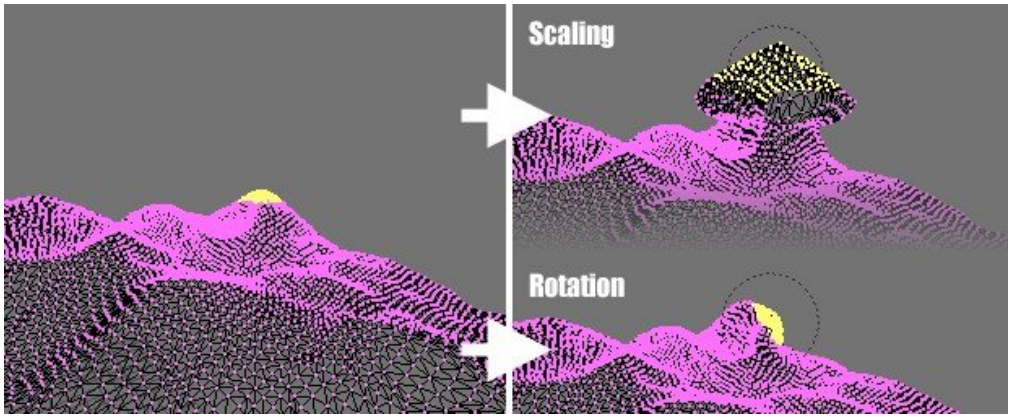
Influence

You can increase or decrease the radius of the proportional editing influence with the mouse wheel WheelUp /WheelDown or PageUp/PageDown respectively. As you change the radius, the points surrounding your selection will adjust their positions accordingly.

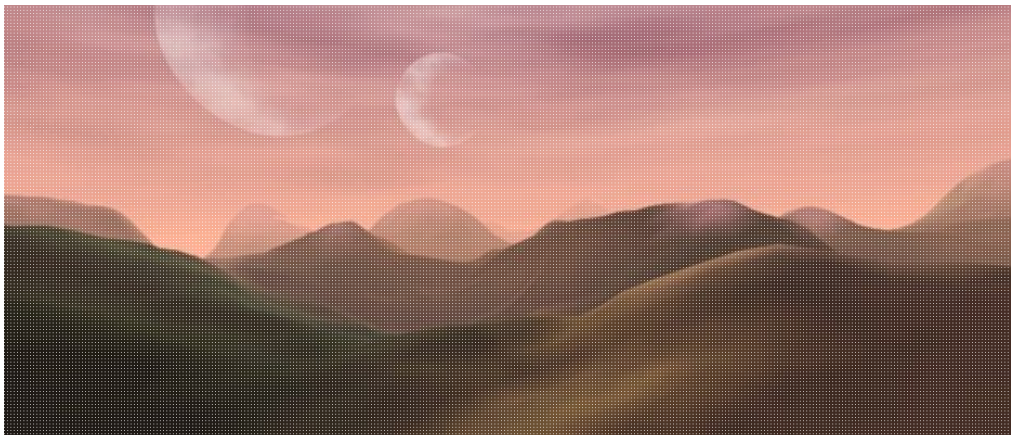
Examples

Switch to a front view (1 NumPad) and activate the grab tool with G. As you drag the point upwards, notice how nearby vertices are dragged along with it. When you are satisfied with the placement, clic LMB to fix the position. If you are not satisfied, cancel the operation and revert your mesh to the way it looked before with RMB (or Esc key).

You can use the proportional editing tool to produce great effects with the scaling (S) and rotation (R) tools, as (*A landscape obtained via Proportional Editing*) shows.



A landscape obtained via Proportional Editing.
Combine these techniques with vertex painting to create fantastic landscapes. (*Final rendered landscape*) shows the results of proportional editing after the application of textures and lighting.



Final rendered landscape.

Transform Orientations

Mode: Object mode and Edit mode

Hotkey: AltSpace



Transform orientations selection menu.
A secondary orientation can be selected with AltSpace or through the Orientation menu in a 3D view header. This orientation can then be used as a transformation constraint.

Axis Locking

Mode: Object mode and Edit mode

Hotkey: XX, YY, ZZ

Pressing the axis locking key twice locks to the user selected transform orientation.

Plane Locking

Mode: Object mode and Edit mode

Hotkey: â § ShiftXâ § ShiftX, â § ShiftYâ § ShiftY, â § ShiftZâ § ShiftZ

The same principle applies to locking movement on a plane. Press the plane locking keys twice to lock to your selected transform orientation.

Custom Orientations

Mode: Object mode and Edit mode

Hotkey: â § ShiftCtrlC

You can define custom transform orientations, using object or mesh elements. Custom transform orientations defined from objects use the local orientation of the object whereas those defined from selected mesh elements (vertices, edges, faces) use the normal orientation of the selection.

A name also needs to be assigned to the new orientation.

Note
When adding new orientations, if the name correspond to an already existing custom orientation, the new orientation will replace the old one.

[Demo Video] (XviD)

[Bits of Blender] has a video tutorial on this topic:

Transform Orientations Panel

The Transform Orientations panel, available from the View menu, can be used to manage transform orientations: selecting the active orientation, adding and removing custom orientations and clearing all custom orientations.



Transform Orientations panel.



Selecting a custom orientation.



Transform Orientations panel with custom orientation selected.

Transform Properties Panel

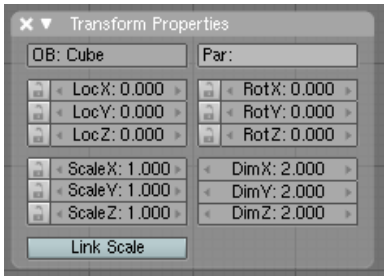
Mode: Edit and Object modes

Hotkey: N

Menu: Object â Transform Properties

The Transform Properties floating panel allows you to manually/numerically control the position, rotation, and other properties of an object, in Object mode. In Edit mode, it mainly allows to enter precise coordinates for a vertex, or median position for a group of vertices (including an edge/face). As each type of object has a different set of options in its Transform Properties panel in Edit mode, see their respective descriptions in the Modeling chapter.

Options in Object mode



Transform Properties panel in Object mode.

OB
: The objectâ← s name.

Par
: The name of the parenting object, if one is assigned. By entering in a name you are assigning a parent object. The name must match an existing object (if it doesnâ← t, the name is erased from the field).

LocX, LocY, LocZ
: The objectâ← s center location in global coordinates.

RotX, RotY, RotZ
: The objectâ← s orientation, relative to the global axes and its own center.

ScaleX, ScaleY, ScaleZ
: The objectâ← s scale, relative to its center, in local coordinates (i.e. the ScaleX value represents the scale along the local X-axis). Each object (cube, sphere, etc.), when created, has a scale of one blender unit in each local direction. To make the object bigger or smaller, you scale it in the desired dimension.

DimX, DimY, DimZ
: The objectâ← s basic dimensions (in blender units) from one outside edge to another, as if measured with a ruler. For multi-faceted surfaces, these fields give the dimensions of the bounding box (aligned with the local axes â← think of a cardboard box just big enough to hold the object).

Link Scale

If this toggle-button is activated the relation (proportion) between the X, Y and Z values in the Scale and the Dim fields is always preserved. Changing one value will change all the others as well with the same multiplication-factor.

Use this panel to either edit or display the objectâ ← s transform properties such as position, rotation and/or scaling and this includes the objectâ ← s name and parent assignment. These fields change the objectâ ← s center and then affects the aspect of all of its *vertices* and faces.

Ipo Note
The values of the location, rotation, and scale can also be affected by an Ipo keyframe, so if there are Ipo keys associated with the object, be sure to reset them after making changes in this panel, or your changes will be lost when the frame is changed (Ipo keys override manually-set properties).

Some fields have extra functionality or features, such as scroll regions. When attempting to edit these types of fields it is easier to use {â § ShiftLMB instead of just LMB. After you have edited a field click outside of the fieldâ ← s edit area or press â µ Enter to confirm the changes. Changes will be reflected in the display window immediately. To cancel, hit Esc. For further descriptions of the other features of an edit field see The Interface section.

Transform Properties Locking

The locking feature of the Location, Rotation and Size fields allows you to control a transform property solely from the properties panel. Once a lock has been activated any other methods used for transformation are blocked. For example, if you locked the LocX field then you canâ ← t use the mouse to translate the object along the global X axis. However, you can still translate it using the LocX edit field. Consider the locking feature as a rigid constraint only changeable from the panel.

To lock a field, click the padlock icon next to the field. The field is unlocked if the icon appears as (), and it is locked if the icon appears as ().

Overview

Each .blend file contains a database. This database contains all scenes, objects, meshes, textures, etc. that are in the file. A file can contain multiple scenes and each scene can contain multiple objects. Objects can contain multiple materials which can contain many textures. It is also possible to create links between different objects.

Mode: All Modes, Any Window

Hotkey: â § ShiftF4 - Data Select Browser

To access the database, press â § ShiftF4 and the window will change to a Data Select Browser window, which lists the Objects in your .blend file. To go up a level, click the breadcrumbs (. .) and then you will see the overall structure of a file: Action, Armature, Brush, Camera, Curve, Group, and so on (including Objects).

LMB  selecting any datablock type, Mesh, for example, will give you a listing of the meshes used in the file, along with how many users there are for each one. For example, if you had a car mesh, and used that car mesh for six cars in a parking lot scene, the Mesh listing would show the Car and then the number 6.

Mode: Data Select Browser

Hotkey: F - Fake User

RMB  selecting certain kinds of datablocks (Materials, Images, Texturesâ ←¦) and pressing F will assign a â ← fake userâ ← to those datablocks. With a fake user in place, Blender will keep those datablocks in the file, even if they have no â ← real userâ ← . Datablocks without a user, real or fake, are not saved in the .blend file. Pressing F again toggles the fake user assignment. Performing this action is the same as clicking the F button next to material/image/â ←¦ names.

Outliner and OOPS Schematic

You can easily inspect the contents of your file by using the Outliner window. This window displays the Blender data system (fully documented here). This window offers two views of the database. The Outliner view allows you to do simple operations on the objects. These operations include selecting, renaming, deleting and linking. The OOPS Schematic (Object-Oriented Programming System) view allows you to easily see how datablocks are linked. You can filter the view by using buttons found in the header.

Users (Sharing)

Many datablocks can be shared among other datablocks - re-use is encouraged. For example, suppose you have a material for one object, named â← `Glossy`â← . You can select a second object, for example, one that does not have a material yet. Instead of clicking ADD NEW for the material, click the little up-down arrow next to the ADD NEW, which brings up a list of existing materials. Select â← `Glossy`â← . Now, these two objects share the same material. You will notice a â← 2â← next to the name of the material, indicating that there are two users (the two objects) for this material. Other common examples include:

- Sharing textures among materials.
- Sharing meshes between objects (â← clonesâ←).
- Sharing Ipo curves between objects, for example to make all the lights dim together.

Fake User

Remember that Blender does not save datablocks that are not linked to anything in the *current* file. If you're building a ".blend" file to serve as a library of things that you intend to link-to from *other* files, you'll need to make sure that they don't accidentally get deleted from the current (the library) file. Do this by giving the datablocks a "fake user," by hitting the F button next to the name of the datablock. This prevents the user count from ever becoming zero: therefore, the datablock will not be deleted. (Blender does not keep track of how many other files link to this one.)

Copying and Linking Objects Between Scenes

Sometimes you may want to link or copy objects between scenes. This is possible by first selecting objects you want to link or copy and then using the Make Links and Make Single User items found in Object menu in the 3D viewport header. Use Make Links to make links between scenes. To make a plain copy, you first make a link and then use Make Single User to make a stand-alone copy of the object in your current scene. Further information on working with scenes can be found here.

Appending or Linking Across Files

The content of one .blend file is easily accessed and put into your current file by using the File â Append function (accessed at any time by â § ShiftF1). To find out more about how to copy or link objects across .blend files, click here.

Proxy Objects

Proxy objects allow you to make (parts of) linked data local. For example, this allows an animator to make a local â← copyâ← of the handler bones of a character, without having the actual rig duplicated. This is especially useful for character animation setups, where you want the entire character to be loaded from an external library, but still permit the animator to work with poses and actions. Another example: you can have a modeler working on the shape (mesh) of a car and another painter working on the materials for that car. The painter cannot alter the shape of the car, but can start working with color schemes for the car. Updates made to the shape of the car are applied automatically to the painterâ← s proxy.

Pack and Unpack Data

Blender has the ability to encapsulate (incorporate) various kinds of data within the .blend file that is normally saved outside of the .blend file. For example, an image texture that is an external `.jpg` file can be put â← insideâ← the .blend file via File â External Data â Pack into .blend file. When the .blend file is saved, a copy of that `.jpg` file is put inside the .blend file. The .blend file can then be copied or emailed anywhere, and the image texture moves with it.

You know that an image texture is packed because you will see a little â← Christmas present gift boxâ← displayed in the header.

Unpack Data

When you have received a packed file, you can File â External Data â Unpack into Files.... You will be presented with the option to create the original directory structure or put the file in the `//` (directory where the .blend file is). Use â← original locationsâ← if you will be modifying the textures and re-packing and exchanging .blend files, so that when you send it back and the originator unpacks, his copies of the textures will be updated.

Datablocks

The base unit for any blender project is the datablock. Examples of datablocks include: meshes, objects, materials and textures. Be it a simple sphere floating over a plane, or a full featured film, a project is defined by the datablocks it contains, the properties set for those datablocks, and the way the datablocks link to each other. Datablocks can reside within as many .blend files as needed for good project organization.

Scene Datablock

Each Scene datablock contains a scene. This type of datablock is the parent of all the others.

Object Datablock

Each Object datablock has the properties of Location, Rotation and Scale, and is the â← meeting placeâ← for other datablocks that define the other properties of that object when they are linked to it. They can determine its nature: mesh, curve, camera, lamp or armature datablocks are a few examples of such datablocks. Others will define the material, the texture, the animationâ←¦ for that object.

ObData Datablocks

These datablocks are always connected to Object datablocks in one way or another.

Curve Datablock

Curve datablocks may contain NURBS curves, circles or surfaces, BÃ©zier curves or circles, or Text objects. It may also be linked to a Material datablock.

Camera Datablock

Camera datablocks contain cameras.

Lattice Datablock

Lattice datablocks contain lattices.

Lamp Datablock

Lamp datablocks contain lamps.

Metaball Datablock

Metaball datablocks contain metaballs.

Mesh Datablock

Mesh datablocks contain meshes. These may contain links to one or more Material datablocks.

Material Datablock

Material datablocks contain materials. They may contain links to Texture datablocks.

Note that Material datablocks can be linked to Object datablocks instead if desired. You can set this in the User Preferences window, below Edit Methods.

Texture Datablock

Texture datablocks contain textures (procedural, imageâ←¦).

Copying and Linking Datablocks

It is possible to copy or link any type of datablock.

Copying and Linking Scene Datablock

To copy a Scene datablock, use Scene list found in the header of User Preferences window. The list is to the right of the menus and window workspace list. Select ADD NEW to make a copy of the current scene. Select Full Copy from the list that opened to make a copy. The current scene will be **fully** copied to the new scene.

Instead of copying **everything**, you can link datablocks by selecting Link Objects, to use the same Object datablocks linked into the two scenes, or Link ObData, to create separated objetcs sharing the same ObData datablocks (meshes, curves, materials, etc.). Note that if you select Link Objects, in fact you copy nearly nothing, as Object datablocks are parent of all ObData datablocks: nearly all modifications (object location, mesh geometry, â←¦) in a scene will be reflected in the other ones linked this way as well. As for Link ObData, it creates a â← realâ← copy of the objects, but not of their child datalocks: the locations, scales and rotations of the objects will be specific to a scene, but neither their forms nor their materials, textures, â←¦, will be (as defined by ObData datablocks).

Copying and Linking Object Datablocks

Full copy
: â § ShiftD is used to make normal copy of the selected objects.
The object and some of itâ← s child datablocks will really be duplicated, the other children are just linked; you can define the attributes to be duplicated in User Preferences â Edit Methods, button group Duplicate with object:.

Linked copy
: AltD makes a linked copy.
All datablocks but the object one are linked.

Copying and Linking other Datablocks

When an ObData datablock is used (linked) by more than one object, a small button with its number of linked objects (users) shows up next to its name, in its settings window (Editing context for meshes, curves, cameras, â←¦, Shading context, Material sub-context for materials, etc.). If you click on it, you create a copy of this datablock for the current object.

Unlinked Datablocks

A datablock can become unlinked. For example a material datablock will be unlinked if the object it was linked to is deleted. If a datablock is unlinked, by default it will be deleted from the .blend file when Blender is closed. To keep an unlinked datablock in a .blend file, click the "F" button to the right of the object's name in the Objects and Links panel.

Scene Management Structure

Scene management and library appending/linking is based on Blenderâ← s Library and Data System, so it is a good idea to read that manual page first if youâ← re not familiar with the basics of that system.

Blender can be used to create something as simple as a single scene or image, or scaled up to an entire movie. A movie is usually comprised of three acts:

1. Introduction-Conflict.
2. Rising Tension.
3. Climax-Resolution.

Each act contains a few scenes, settings where the action happens. Each scene is shot on a set, stage or location. Each is set with props and backdrops. The scene is a set of action sequences where the actors act (hopefully convincingly). Each sequence, or shot, usually lasts a few seconds.

Sequence shot
Sometimes, a single shot lasts several minutes: its a â← sequence shotâ← , which might even be a complete scene on its own. Technique hard to master if you donâ← t want your audience to fall asleep!

A single Blender file is organized and set up to be able to contain an entire movie. Each .blend file can contain multiple scenes. A scene is a bunch of objects, organized into layers. As you progress through the creative process, you use a set of window screen layouts specifically designed to help you work efficiently through the creative process: model the objects and create the props, clothe the actors and dress the set (assign materials), define the action (animation), render the video, and produce the movie. You can tailor these screen layouts, and create custom layouts, to match your working preferences.

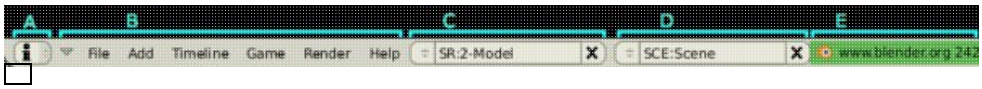
Planning Your Timeline

Shots within a scene are accomplished by moving a camera and/or actors through the scene for a few seconds. Time in Blender is measured in frames, and typical video has 25 or 30 frames per second (fps), and film is 24 fps. For a five-second shot then, you allocate up to 150 frames for that shot (30 fps Ã 5 seconds). Giving yourself some wiggle-room, shot 2 would start at frame 250 and go from there. A one-minute movie set in a single scene for North America video broadcast (NTSC standard) would have a timeline that goes up to 1800 final frames, and may be laid out over the course of 2500 frames. This timeline allows for cutting out 700 frames, picking the best 1800 frames (30 fps Ã 60 seconds = 1800 frames) less transition time.

Multiple Cameras
You can have multiple cameras in a scene, used for different shots, and select which one is active when rendering the shot.

Current Screen Layout and Scene

Scenes are a way to organize your work. Scenes can share objects, but can, for example, differ from one another in their rendered resolution or their camera view. The current window layout and scene are shown in the User Preferences window header, usually displayed at the top of your screen:



User Preferences window header. **A**) Window type icon,
B) Menu, **C**) Screen Layouts, **D**) Scenes,
E) Version of Blender currently running (click the Blender icon to the left to show splash screen).

Loading the UI with â← Fileâ← â â← Openâ←

Inside each .blend file, Blender saves the user interface layout - the arrangement of screen layouts when the file is saved. By default, this saved UI is loaded, over-riding any user defaults or the current screen layout. However, you can work on a blend file using your current UI settings by ignoring the UI settings saved in the file. This is done by restarting Blender or resetting it with (File â New, or CtrlX), and opening the file browser with (File â Open..., or F1). Turn off the Load UI button in the file browser header, and then open the file. This way, Blender will not disturb your current screen layout when it loads the new file.

Working with Scenes

Select a scene to work on by clicking the up-down arrow next to the scene name. Scenes and the objects they contain are generally specific to the project you are working on. However, they too can be saved in their current state to be re-used by pressing CtrlU. They will then appear the next time Blender starts or when the user selects File â New (CtrlX).

Blender comes with one default scene, which contains a camera, a lamp, and a box.

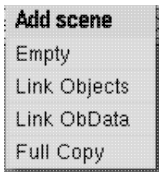
Adding a Scene

You can make a full copy of the current scene, start over with a blank slate, or create a scene that has links back to the current scene; objects will show up in the new scene, but will actually exist in the old one. Use this linking feature when, for example, the original scene contains the set, and the new scene is to contain the actors or props.

Starting Over
If you start with a new scene, be sure to add a camera and lights first!

Scenes are listed alphabetically in the drop-down list. If you want them to appear in a different order, start them with a numerical ordinal, like â← 1-â← . The internal reference for a scene is the three-letter abbreviation â← SCEâ← .

To add a scene, click on the scene list button, and select Add New. While you are adding a new scene, you have these options:



Add scene popup menu.

Empty
: Create a completely empty scene.

Link Objects
: All objects are linked to the new scene. The layer and selection flags of the objects can be configured differently for each scene.

Link ObData
: Duplicates objects only. ObData linked to the objects, e.g. mesh and curve, are not duplicated.

Full Copy
: Everything is duplicated.

Usually, for your first scene, you make a full copy of the default. Alternatively, you can just start with the default, and start editing the cube that is usually hanging around waiting for you to do creative things.

Naming a Scene

By â § ShiftLMB -clicking on the scene name (usually â← `Scene.001`â←), you can change the name of the scene. For example, â← `BoyMeetsGirl`â← is usually the first of three acts.

You then proceed to model the props and objects in the scene using the 2-Model window layout.

Linking to a Scene

You can, at any moment, link any object from one scene to another. Just open the scene where these objects are, use CtrlL â To Scene..., and chose the scene where you want your objects to appear. Those will be linked to the original objects; to make them single user (independent, unlinkedâ←¦) in a given scene go to that scene, select them and use U. You will be presented with a few options that allow you to free up the datablocks (Object, Material, Textureâ←¦) that you want.

Removing a scene from the file

You can delete the current scene by clicking the X next to the name.

Outliner window

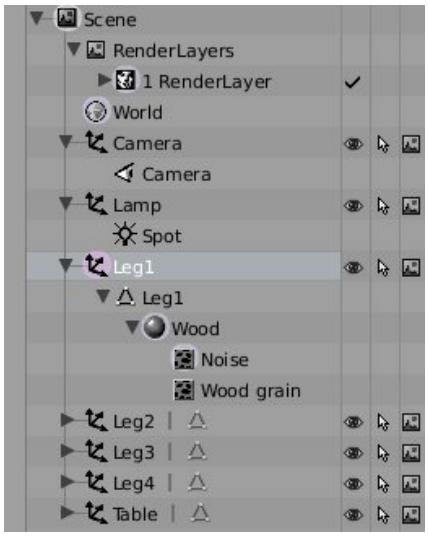
Description

The Outliner window is used for easily navigating a complex scene. There are two views, the Outliner view and the Oops Schematic view. The Oops Schematic and Outliner give you a 2D representation of your complicated 3D world. Use these views to find things in your scene.

For example, suppose you sneeze while moving an object; your mouse flies off your desk (gesundheit!) and the object is hurled somewhere off screen into space. Simply use the schematic/outliner to find it; select it, and move back to your 3D window to snap it back to your cursor (â § ShiftS â Selection -> Cursor).

Another more practical example is to evaluate the impact of a change on related datablocks. Suppose you are looking at your `TableTop` object, and it doesnâ ← t look right, the `Wood` material doesnâ ← t look right; you want it to look more like mahogany. Since the same material can be used by many meshes, youâ ← re not sure how many things will change color when you change the material. Using the Oops Schematic, you could find that material and trace the links that it has to every mesh in your scene.

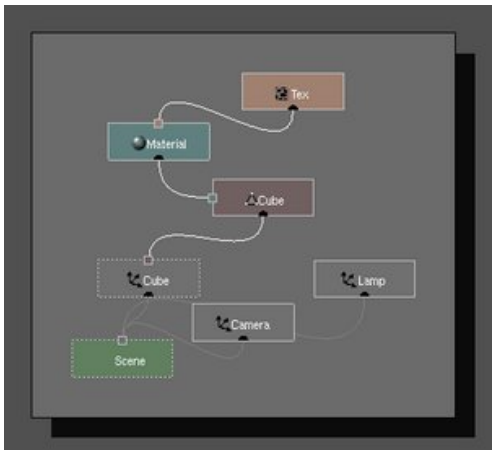
Outliner view



The Outliner window in Outliner mode.
The Outliner is a kind of list that organizes related things to each other. In the outliner, you can:

- View the data in the scene.
- Select and deselect objects in the scene.
- Hide or show an object in the scene.
- Enable or disable selection (to make an object â ← unselectableâ ← in the 3D Views).
- Enable or disable the rendering of an object.
- Select data like materials and textures directly (they show up automatically in the Buttons window!).
- Delete objects from the scene.
- Unlink data (equivalent to pressing the X button next to the name of a datablock).
- Easily select which render layer to render.
- Easily select which render pass to render (for example, you can choose to render just the Specular pass).

Oops Schematic view

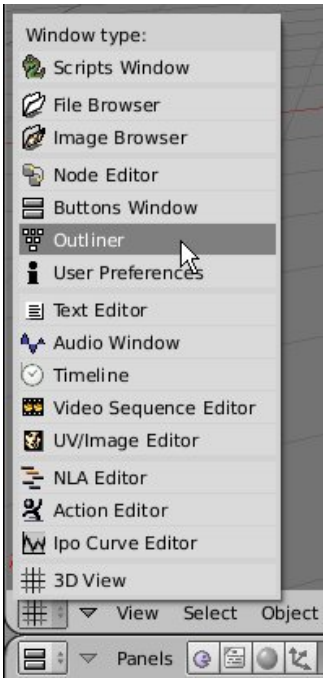


The Oops Schematic view in the Outliner window.
The Oops Schematic is a kind of picture that shows how things are linked together. â← **Oops**â← is a highly geeky term for â← Object-Oriented Programming Systemâ← . In the Oops view, you can:

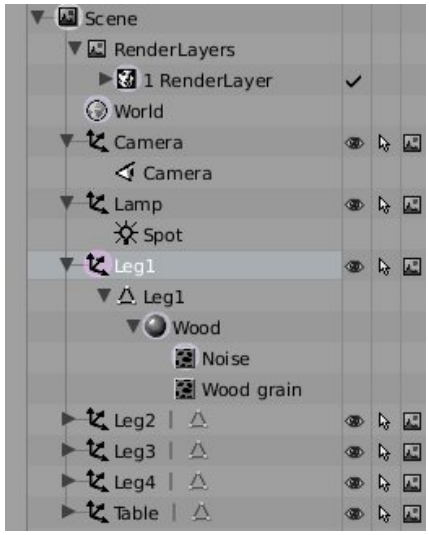
- Look at relationships *between* objects (for example, which objects use the same texture).

The main difference is that the Oops Schematic shows you *all available* things (datablocks) in your blend file, organized by type of thing: scenes at the bottom, objects in the middle, materials toward the top. The Outliner shows you things *in use* within your blend file, organized by parent object with their children as indents.

Selecting the Outliner Window Type



Turn a window into Outliner type.



The Outliner window in Outliner mode.

Choose a window and click on its current Window type button (left-most icon in its header), and select Outliner.

Switch between the Outliner view and the Oops Schematic view using the menu item View â Show Oops Schematic or View â Show Outliner.

Window size
Choose or arrange a window size that suits the view you are going to work with. The Oops Schematic needs a wide window, and the Outliner needs a tall, narrow window.

Using the Outliner view

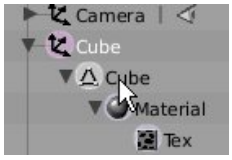
Each row in the Outliner view shows a datablock. You can click the down-arrow to the left of a name to expand the current datablock and see what other datablocks it contains.

You can select datablocks in the Outliner, but this wonâ ← t necessarily select the datablock in the scene. To select the datablock in the scene, you have to activate it.

Selecting and activating

Single selection doesnâ ← t require any pre-selection: just work directly with LMB (and/or RMB - contextual menu, see bellow) *inside* the name/icon area.

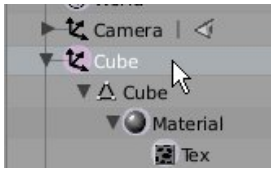
When you select an object in the list this way, it is selected and becomes the active object in all other 3D Views. Use this feature to find objects in your 3D View, select them in the Outliner, then zoom to them with . NumPad or if you don't have a numpad, snap and center your cursor on them via â § ShiftS â Cursor -> Selection, and then C.



Click LMB on the mesh data of the cube to activate Edit mode.

Activating a datablock

Activate the datablock with LMB on the *icon* or the *name* of the datablock. Activating the datablock will automatically switch to the relevant mode or Buttons context. For example, activating the mesh data of the cube will select the cube and enter Edit mode (see right). Another example is that activating the material datablock for the cube will show this material in the Shading context, Material sub-context of the Buttons window.



Toggling pre-selection of a datablock.

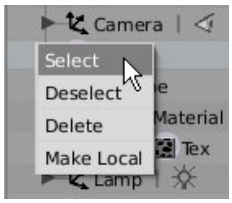
Toggle pre-selection of a group of datablocks

Useful when you want to select/deselect a whole bunch of datablocks. For this you must prepare the selection using, to your liking:

- ◊ RMB or LMB,
- ◊ â § ShiftRMB or â § ShiftLMB,
- ◊ RMB and drag or LMB and drag,

all *outside* the name/icon area. Those pre-selected have their line in a lighter color.

You then can (de)select them with a RMB *on* the name/icon area, which brings on a context menu (see bellow).



Context menu for the Cube object.

Context menu
: Show the context menu for a datablock with RMB on the icon or name. Depending on the type of the pre-selected datablock(s), you will have all or part of the following options:
 ◊ Select.
 ◊ Deselect.
 ◊ Delete.
 ◊ Unlink â← To unlink a datablock from its â← ownerâ← (e.g., a material from its mesh).
 ◊ Make Local â← To create a â← localâ← duplicate of this datablock.
 Note: some datablock types will not have a context menu at all!

Deleting a datablock
: Use X to delete the selected datablock(s).

Expanding one level
: Use + NumPad to expand one level down in the tree-list.

Collapsing one level
: Use - NumPad to collapse one level up in the tree-list.

Expanding/collapsing everything
: Use A to expand/collapse all levels of the tree-list.

Toggling object-level restrictions



Icons for toggling.
The three following options, in the right side of the Outliner window, are only available for objects:

Visibility
: Toggle visibility by clicking the â← eyeâ← icon for the object on the right-hand side of the Outliner. Useful for complex scenes when you donâ← t want to assign the object to another layer. This will only work on visible layers - an object on an invisible layer will still be invisible regardless of what the Outliner says. V will toggle this property for any objects that are pre-selected in the Outliner.

Selectability
: Toggle selectability by clicking the â← arrowâ← icon. This is useful for if you have placed something in the scene and donâ← t want to accidentally select it when working on something else. S will toggle this property for any objects that are pre-selected in the Outliner.

Rendering
: Toggle rendering by clicking the â← cameraâ← icon. This will still keep the object visible in the scene, but it will be ignored by the renderer. R will toggle this property for any objects that are pre-selected in the Outliner.

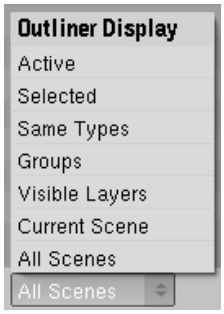
Searching

You can search the file for datablocks, either by using the Search menu in the header of the Outliner, or by using one of the following hotkeys:

- F - Find.
- CtrlF - Find (case sensitive).
- AltF - Find complete.
- CtrlAltF - Find complete (case sensitive).
- â § ShiftF - Find again.

Matching datablocks will be automatically selected.

Filtering the display

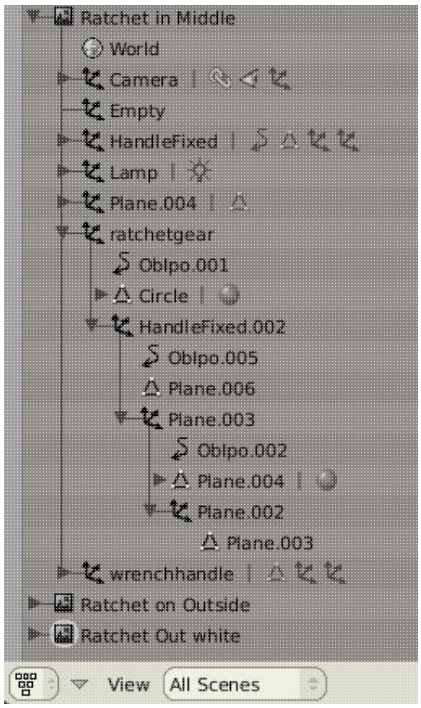


Outliner Display dropdown.
The window header has a field to let you select what the outliner should show in the outline. By default, the outliner shows All Scenes. You can select to show only the current scene, datablocks that have been selected, objects that are on currently selected layers, etc. These selects are to help you *narrow the list* of objects so that you can find things quickly and easily.

- All Scenes - Shows *everything* the outliner can display (in all scenes, all layers, etc.)
- Current Scene - Shows everything in the current scene.
- Visible Layers - Shows everything on the visible (currently selected) layers in the current scene. Use the layers buttons to make objects on a layer visible in the 3D window.
- Groups - Lists only Groups and their members.
- Same Types - Lists only those objects in the current scene that are of the same types as those selected in the 3d window.
- Selected - Lists only the object(s) currently selected in the 3D window. You can select multiple objects by â § ShiftRMB -clicking.

- Active - Lists only the active (often last selected) object.

Example



The Outliner window in list mode.
The outline example shows that the .blend file has three scenes: â← `Ratchet in Middle`â← , â← `Ratchet on Outside`â← , and â← `Ratchet Out White`â← . By clicking on the little arrow to the left of the name, the outline is expanded one level. This was done for the â← `Ratchet in Middle`â← scene. As you can see, this scene has some â← `World`â← material settings, a â← `Camera`â← , an â← `Empty`â← , a â← `HandelFixed`â← objectâ←¦ All objects that were added to the scene.

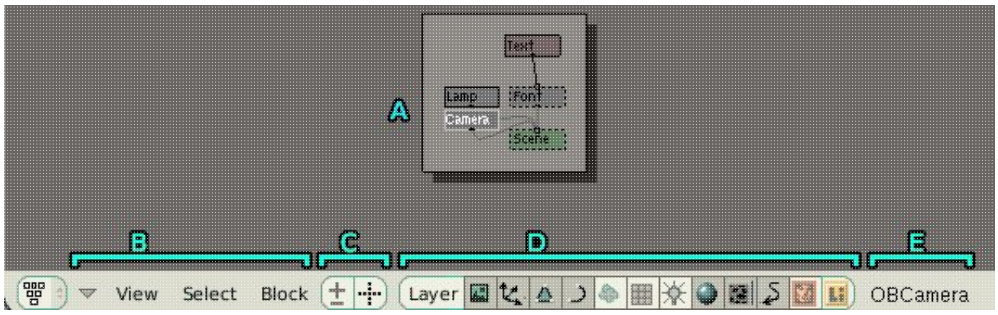
By clicking the arrow next to â← `ratchetgear`, we can see that it has some motion described by the â← `ObIpo.001`â← curve; that it was based on a â← `Circle`â← mesh, and that it is the parent of â← `HandleFixed.002`â← , which is in turn the parent of â← `Plane.003`â← , and so on.

The neat thing is: if you select any of these datablocks here, they will be selected in the 3D (or Buttons) window as well as far as this is possible. Pressing * NumPad with your mouse cursor in any 3D Window will center and align the view to that object. Very handy. Also, pressing X will delete it, as well as all the other hotkeys that operate on the currently selected object.

Using the Oops Schematic

Layout of the Oops Schematic

In this view, the window has a clear background that, by default, shows the Oops Schematic and a header:



Outliner window in Oops Schematic mode.
A) OOPS Schematic, **B**) Menu, **C**) Zoom, **D**) Display Select, **E**) Currently selected datablock.

The Oops Schematic window & header have the following areas:

A) The schematic picture.
B) Menus with the basic functions: View, Select, and Block.
C) A zoom control that allows you to focus on a certain area of the schematic.
D) Visible Select - a number of buttons that toggle what kinds of datablocks are displayed in the schematic.
E) The name of the currently selected datablock. The datablock is also highlighted in the Oops schematic (**A**).

Making sense of the Oops Schematic

The schematic is a sort of map that shows the connections between datablocks. Each datablock is shown as a colored box. Boxes (datablocks) are connected by lines. Common types of connections between datablocks are:

Parents
: One datablock, letâ← s say an object called â← `TableTop`â← , is held up by four other objects â← `leg.001`â← , â← `leg.002`â← , etc. The `TableTop` would be the parent of each of the legs, so that as the table top moves, the legs move as well. In the schematic, four lines would be shown going from the `TableTop` to each of the Legs.

Material Use
: Datablocks can share the same material. In our table example, the `TableTop` and each of the legs might share the same material, â← `Wood`â← , so that they all look the same. In the schematic, there would be a box called â← `Wood`â← with five lines connecting it to each of the mesh datablocks `TableTop`, `Leg.001`, `Leg.002`, `Leg.003` and `Leg.004`.

The schematic uses different colored boxes for each type of datablocks, to help you visually distinguish between them:

- Grey-green for scenes.

- Grey for objects.
- Olive green for groups.
- Purple-grey for meshes, *curves, surfaces, metas, texts.*
- Beige for armatures.
- Purple for lattices.
- Greyish yellow for lamps.
- Pink for cameras.
- Sea green for materials.
- Salmon pink for textures.
- Orange for images.
- Blue for Ipos.
- The datablocks linked from a library file have a small colored square.
- The non-linked datablocks became uncolored (transparent background).

The Oops Schematic Header

View

Handy functions include switching between the schematic and outliner view. Also, you can change the size of the boxes, so more can fit in the window.

Select

Key functions include finding users and links between connected boxes.

Block

Scales (S) the distance between multiple selected datablocks, and grabs/moves (G) a datablock or set of selected datablocks around the schematic - very useful for arranging and organizing your schematic.

Zoom controls

As you can imagine, depending on what you have selected and your scene complexity, these schematics can start looking like the piping diagram for a nuclear power plant. The schematic header provides two buttons to help you zoom in.

- ◊  button: hold down LMB over this button and move your mouse up and down (forward and backward) to zoom your view in and out.
- ◊  button: click on this button to start a border select. Select a region in the window, and your window view will be zoomed to that region.

Standard Window Controls

Like in any Blender window, you can also zoom in/out with WheelUp /WheelDown , or click and drag CtrlMMB . You can pan it to, with a click and drag MMB for exemple.

Visible Select

The series of icons in the header allow you to select what type(s) of datablocks (see this page for more details about datablocks) are visible in the schematic. They are, left to right:

- ◊  - Only show the datablocks from the shown layers.
- ◊  Scenes - Your stage, a set, where action occurs.
- ◊  Objects - All Objects datablocks.
- ◊  Meshes - The main things you model, not to be confused with objects. E.g. one Mesh can be used in multiple objects and is displayed accordingly in the schematic.
- ◊  Curves, Surfaces, Texts - The â← realâ← data of parametric objects.

◊ Metaballs - Mathematically calculated meshes that can mush together, as a fluid in zero-G.
◊ Lattices - Deformation grids.
◊ Lamps - All types of lights.
◊ Materials - Colors, shadingâ←¦
◊ Textures - Color maps or gradients used commonly in materials and other places.
◊ Ipos - All animated settings (positions, rotations, colors, â←¦).
◊ Images - All images loaded in Blender, as a texture, in the VSE, â←¦
◊ Groups - The groups of objects.
◊ Libraries - Collections of objects.
◊ Cameras - The â← eyesâ← of the Blender.
◊ Armatures - The skeletons of your characters.

Linked Libraries Overview

Blender is able to â← reach inâ← to other .blend files and pull in whatever you want. In this way, Blender supports reuse of your graphical models. For example, if you have a library .blend file that has a really neat material used in it, you can, from your current .blend file, Append that material into your current .blend file. This saves you from manually re-creating all the different settings.

General Procedure

Mode: All Modes

Hotkey: â § ShiftF1

Menu: File â Append or Link

The main menu in Blender is located in the User Preferences window (by default the header located at the top of your screen). From that menu, all you have to do is use File â Append or Link, or press â § ShiftF1 in your active window. The active window will change to a File Browser (the Window type icon looks like a manila folder) selector window. Use this window to navigate your hard drive and network-mapped drives through folders and subfolders to find the .blend file that has the object you want to reuse. When you click on a .blend file (indicated by the orange square box next to its name), Blender will go into that file and show you the list of datablock types within it: Scenes, Objects, Materials, Textures, Meshes, etc. Clicking on any one of them will display the specific instances of that type.

Folder and File Organization

We suggest creating a folder called `/lib` or `/library`. Under that library, create a set of folders for each kind of thing you might want to access and re-use later on, such as `materials`, `textures` and `meshes`. Create subfolders under each of those as your library grows. For example, under the `meshes` folder, you might want to create folders for `people`, `spaceships`, `furniture`, `buildings`, etc. Then, when you have a .blend file that contains a chair mesh, for example, all you have to do is copy that file into the `furniture` folder.

Appending library objects into your current project

The following procedure appends an object with all its linked data, such as mesh data, materials, textures, â←¦, to the current .blend file.

1. Select File â Append or Link.
2. Locate and select the file that contains the object you want to append (often a â← libraryâ← file).
3. Navigate to the Object section of the file.
4. Select one object from the list using LMB, multiple objects via RMB, and/or a range of objects by dragging RMB.
5. Repeat the above for each kind of object you wish to append or link. Parents and armatures (all modifier objects) must be selected separately.
6. Set desired options that are shown in the header (At Cursor, Active Layer, â←¦).
7. LMB on Load Library or press â µ Enter or MMB directly on the data to append.

Of course, you can append or link many other things besides objects: all the ObData - cameras, curves, groups, lamps, materials, meshes, etc. - and even **an entire scene**â←¦ Note that there is a **big** difference

between adding the object and the object data, such as mesh. If you append a Mesh datablock, you are only bringing in the data about that particular instance of mesh, and not an actual object instance of the mesh that you can see.

In the File Browser window header, use Append (button enabled by default) if you want to make a local independent copy of the object inside your file. Select Link if you want a dynamic link made to the source file; if anyone changes the object in the source file, your current file will be updated the next time you open it.

Click Load Library to append or link the object into your current .blend file.

Some more loading option buttons (in the File Browser header) include:

AutoSel
: When an object is loaded, it is not active or selected; it just plops into your .blend file. Often, right after loading, you will want to do something with it, like scale it or move it. Enable this button and the imported object will be selected, just as if you magically RMB -clicked on it. This button saves the step of finding the object and selecting it.

Active Layer
: Blender has 20 layers to divide up a large scene, and each object resides on some layer(s). By default, an object is loaded into your file directly into the layer(s) it resides on in the source file. To only load the object to the current active layer that you are working on, enable this button.

At Cursor
: By default, an object is loaded into your file at the location it is at in the source file. To reposition the object to your cursor when it loads, enable this button.



Finding What was Loaded:

If the loaded object is not visible, consider using At Cursor or AutoSel. If you use AutoSel, remember there are Snap tools to put your cursor on the object (â § ShiftS4 (Cursor -> Selection)), and Center your view on it (C (View â Align View â Center View to Cursor)). Note that these tools do not work if the object is on an unselected layer, since objects on unselected layers are invisible.

Reusing Objects (Meshes, Curves, Cameras, Lights, â ←¦)

Letâ ← s suppose you created a wheel in one .blend file and want to reuse it for your current project. The physical model of the wheel would be a mesh, and probably comprised of a tire and rim. Hopefully you named this mesh something reasonable, like, oh, I donâ ← t know, â ← `Wheel`â ← . The wheel may be colored and thus have some materials assigned to it (like rubber and chrome).

Once you navigate to the file, select the â ← `Wheel`â ← (in the Objects datablocks) and it will be imported into your current file. You can import a copy of it, or merely link to it.



Linking:

If you link to it, and later modify it in the source file, it will be shown â ← as-isâ ← (modified) in your current file the next time you open it up.

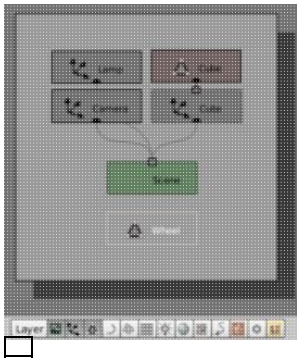
Other artists have released their models to the public domain, and friends may share models simply by posting or emailing their .blend files to each other. Keeping these files, as well as your past projects, in a `Download`

directory on your PC/server will save you from ever having to reinvent the wheel.

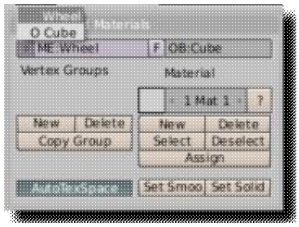
When selected, linked objects are outlined in Cyan. Normal selected objects are outlined in pink.

Notice that you cannot move a linked object! It resides at the same position it has in the source file. To move/scale/rotate the object, turn it into a proxy.

Using Appended/Linked Mesh Data:



When Appending or Linking certain resources such as mesh data, it may not be instantly visible in the 3D Viewport. This is because the data has been loaded into Blender but has not been assigned to an object, which would allow it to be seen. You can verify this by looking in the Outliner window and switching it to OOPS Schematic view (you may need to have the Displays Scene datablock button selected in its header). In the OOPS Schematic picture you can see that â← `Wheel`â← is not linked to any object.



To allow the newly loaded `Wheel` mesh to be assigned to an object, either select a currently visible object or create a new object (such as a cube), then go to the Link and Materials panel and select the `Wheel` mesh from the mesh drop down panel, at that point you should see it, because it has been assigned to an object.

If instead of Appending/Linking to a mesh you instead load the object into Blender, it should be instantly displayed in the 3D Viewport without having to associate an object with the mesh (as it is already done!).

Reusing Material/Texture Settings



Material preview in Image Browser.
Some materials, like glass or chrome, can be very tricky to get â← just rightâ← . The Blender Foundation has released, for example, a Materials CD, which is available for free to download from their site. Using the .blend files on that CD, you can import common materials, like glass, chrome, wood and bananas. This feature saves you a lot of time, as it often means you donâ← t have to be fiddling with all the little buttons and sliders just to re-create a material. I call out the Banana material because it is a great example of using simple procedural materials with a ColorRamp, and a procedural texture, to give a very realistic look. When you navigate to the file, and select Materials, the browser will show you a sphere sample of that material to help you visualize the texture that goes with the name. For more information on using the Image Browser, see the release notes.

Blender Extension: Library
There is also a fantastic Python script called Blender Library that over-arches all of your files and allows you to construct a master library. This script displays a preview and helps you organize your Blender work. Highly recommended; search www.blendernation.com for â← Blender Libraryâ← , it is also stored on the Blender Wiki Scripts section.

Reusing Node Layouts

To reuse noodles (node layouts), open the original (source) file and create a Group for the set of nodes that you think you want to reuse. When you want to import that node group into your current file, select File â Append or Link from the User Preferences window header (or hit F1), and navigate to the file. When you dive into the file, there will be a NodeTree option. Click it and the list of node groups in that file will be listed. LMB -click the one you want and then Load Library.

[Verse]
Verse is an amazing OpenSource collaboration tool that integrates with Blender. Verse enables multiple people to work on, link, and share objects and modifications in Blender files in real time.

Proxy Objects

A proxy is a legal stand-in or substitute for the real thing. In Blender, when you make a linked copy (described above), you cannot edit the object; all you have is a link to it. You cannot add to it or change it, because its source is in another file that is not open.

When working in a team environment, you may want more flexibility. For example, if modeling a car, you may have one person working on the shape of the car (its mesh), but another working on available color schemes (its materials). In this case, you want to grant the painter a Proxy of the object and allow him/her to modify the material settings. More commonly, you will have a character being animated by a team of

animators; they can define poses, but cannot change the characterâ← s colors or armature, only use what is defined by the master rigger.

The important aspect of a proxy object is that it allows you to edit data locally, but also allows specific data to be kept restricted. Data thatâ← s defined as restricted will always be restored from the library (typically on file reading or undo/redo steps). This restriction is defined in the referenced library itself, which means that only the library files can define whatâ← s allowed to change locally.

For poses, you can control this by indicating bone layers as being restricted. A restricted layer is shown with a black dot in it. Use CtrlLMB on a button to restrict or unrestrict that layer.

Mode: Object Mode

Hotkey: CtrlAltP

To make a proxy object for yourself, establish a link to the source object as described above. With that linked copy selected (RMB) and in view (you can see it in the 3D View), press CtrlAltP and confirm the Make Proxy dialog. The object will be named with the original name plus a â← `_proxy`â← suffix. You may now move and modify the proxy. When selected, it will look like a local object (outlined in pink).

You can then edit unrestricted data. For most objects, this includes the location and rotation. You can also animate the objectâ← s location using Ipo curves. For mesh objects, the shape of the mesh is restricted, so you cannot define shape keys. When you reload your file, Blender will refresh your file with any changes made to the original restricted data, but will not reset your changes (unless the owner has).

Modeling in Blender

As you have seen in the Quick Start chapter, the creation of a 3D scene needs at least three key things: Models, Materials and Lights. In this Part we will delve deeper into the first of these issues: modelling. Modeling is the art and science of creating a surface that mimics the shape of a real-world object or fits your imagination of abstract objects.

Objects come in many forms, shapes and sizes, so Blender has many different tools available to help you make your model quickly and efficiently:

Objects
: Working with objects as a whole

Meshes
: Working with the mesh that defines the shape of an object

Curves
: Using Curves to model and control objects

Surfaces
: Modeling a NURBS surface

Text
: Textual tools for putting words in 3D space

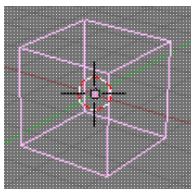
Meta Objects
: Globs and Globules

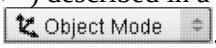
Duplications
: Duplicating Objects

Modeling Scripts
: Since Blender functionality is extensible via Python, there are a number of very useful scripts that assist you in modelling.

Many people use "box modelling" which starts with a basic cube, and proceeds with extruding faces and moving vertices to create a larger, more complicated mesh. For flat objects, like walls and table tops, you can use "curve modelling" which defines the outline using bezier or Nurbs curves, and then extrudes it to the desired thickness. Either method is fully supported in Blender using its modelling tools.

Object Mode



Selected object.
The geometry of a scene is constructed from one or more Objects: For example Lamps, Curves, Surfaces, Cameras, Meshes, and the basic objects (â← primitivesâ←) described in â← Mesh Primitivesâ← . Each object can be moved, rotated and scaled in Object Mode, 

. For making other changes to the geometry of objects, you must use Edit mode.

Once youâ← ve added a basic object (see â← Mesh Primitivesâ←), you are automatically switched into Edit Mode if the Object is a Mesh, a Curve or a Surface. You can switch back to Object Mode by pressing â Tab. The objectâ← s wireframe, if any, should now appear pink, meaning that the object is now selected and active, as shown in (*Selected object*).

Erase

Mode: Edit or Object mode

Hotkey: X or Del

Menu: Object â Delete

Erases or deletes selected objects.

Join

Mode: Object mode

Hotkey: CtrlJ

Menu: Object â Join Objects

Joins all selected objects of the same type to one single object whose center point is obtained from the previously *active* object. Performing a join is equivalent to adding new objects while in Edit mode.

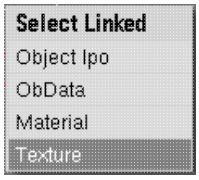
Select Links

Mode: Object mode

Hotkey: â § ShiftL

Menu: Select â Select Linked

Select all objects sharing a link with the active one. You can select objects sharing an Ipo, data, material, or texture link, as shown in (*Selecting links*).



Selecting links.

- Object Ipo: Selects objects that share Ipo information.
- ObData: Selects objects that share *datablock* information (a same mesh, a same curve, a same materialâ←¦).
- Material: Selects objects that share material information.
- Texture: Selects objects that share texture information.

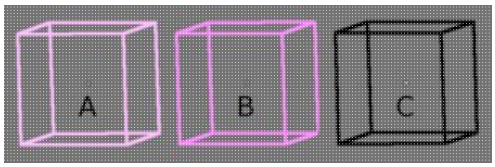
Introduction

Selection, in almost any program, determines which elements will be the target of our actions. As such, the more adapted the selection tool is to the action intended the better. Tools and functions are in a great number in Blender and so are its selection methods.

What follows is a short description of the concepts and selection tools which are available in Object mode.

Selections and the Active Object

Blender distinguishes between two different states of selection:



A) Selected Active Object,
B) Selected Object,
C) Unselected Object.
Outlines have been thickened to make them easier to distinguish.

- In Object mode the last (de)selected item is called the â← Active Objectâ← and is outlined in pink (the others are purple). There is exactly one active object at any time (even when nothing is selected!).

 Many actions in Blender use the active object as a reference, for example the boolean tools or linking operations. If you already have a selection and need to make a different object the active one, simply re-select it with â § ShiftRMB .

- All other selected objects are just that, selected. You can select any number of objects.

Some bugs?
There seems to be display bugs sometimes, like e.g. under (2.49b/linux/ATI Radeon HD 3200), the selected/active objects outlines only show in wireframe viewing mode.

Point Selection

The simplest form of object *selection* consists of using RMB on it.

To *add to the selection*, use â § ShiftRMB on more objects.

If the *objects are overlapping* in the view, you can use AltRMB to get a list of possible choices.

If you want *to add to a selection* this way then the shortcut becomes â § ShiftAltRMB .

To *activate an object* that is already selected, click â § ShiftRMB on it.

To *deselect* an active object, click â § ShiftRMB one time - and hence two clicks if the object isnâ← t active.

Rectangular or Border Select

Mode: Object mode

Hotkey: B

Menu: Select â Border Select

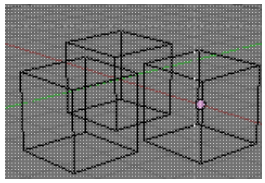
Description

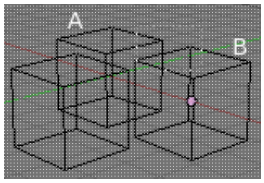
With Border Select you draw a rectangle while holding down LMB. Any object that lies even partially within this rectangle becomes selected.

For deselecting objects, use either MMB or RMB.

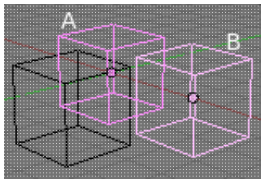
Example

In (*Start*), Border Select has been activated and is indicated by showing a dotted cross-hair cursor. In (*Selecting*), the *selection region* is being chosen by drawing a rectangle with the LMB. The rectangle is only covering cubes â← Aâ← and â← Bâ← . Finally, by releasing LMB the selection is complete; see (*Complete*).


Start.



Selecting.



Complete.

Notice in (*Complete*) the bright color of selected cube â← Bâ← . This means it is the â← active objectâ← , the last selected object prior to using the Border Select tool.

Hints

Border select adds to the previous selection, so in order to select only the contents of the rectangle, deselect all with A first.

Lasso Select

Mode: Object mode

Hotkey: CtrlLMB

Menu: no entry in the menu

Description

Lasso select is used by drawing a dotted line around the pivot point of the objects, in Object mode.

Usage

While holding Ctrl down, you simply have to draw around the pivot point of each object you want to select with LMB.

Lasso select adds to the previous selection. For deselection, use Ctrlâ § ShiftLMB .

Menu Selection

The selection methods described above are the most common. There are also many more options accessible through the Select menu of the 3D view, or the Select option of the â← SpaceBarâ← menu.

Each is more adapted to certain operations.

Select Grouped

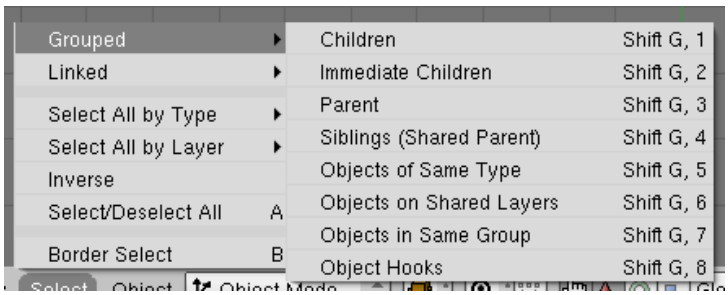
Mode: Object mode

Hotkey: â § ShiftG

Menu: Select â Grouped

Description

There are two ways to organize the objects in relation to one another. The first one is parenting, and the second is simple grouping. We can use these relationships to our advantage by selecting members of respective families or groups.



Options for parented and grouped objects.

Options

Select â Grouped in Object mode uses the active object as a basis to select all others.

Available options are:

Children
: Selects all children of the active object recursively.

Immediate Children
: Selects all direct children of the active object.

Parent
: Selects the parent of this object if it has one.

Siblings (Shared Parents)
: Select objects that have the same parent as the active object. This can also be used to select all root level objects (objects with no parents).

Objects of Same Type
: Select objects that are the same type as the active one.

Objects on Shared Layers
: Objects that have at least 1 shared layer.

Objects in Same Group
: Objects that are part of a group (rendered green with the default theme) will be selected if they are in one of the groups that the active object is in.

Object Hooks
: Every hook that belongs to the active object.

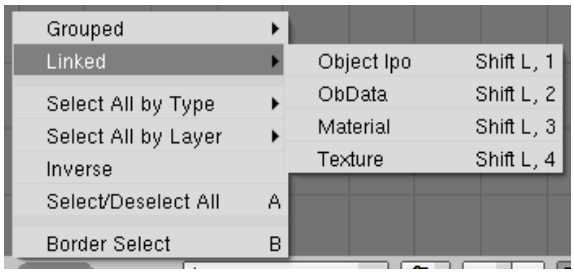
Select linked

Mode: Object mode

Hotkey: â § ShiftL

Menu: Select â Linked

Description



Options for objects which share a datablock.
Selects all objects which share a common datablock with the active object.

Options

Select â Linked in Object mode uses the active object as a basis to select all others.

Available options are:

Object Ipo
: Selects every object that is linked to the same Ipo datablock of the Object type. Any other type like Constraint, Pose, â←¦, wonâ← t work.

ObData
: Selects every object that is linked to the same ObData, i.e. the datablock that specifies the type (mesh, curve, etc.) and the built (constitutive elements like vertices, control vertices, and where they are in space) of the object.

Material
: Selects every object that linked to the same material datablock.

Texture
: Selects every object that linked to the same texture datablock.

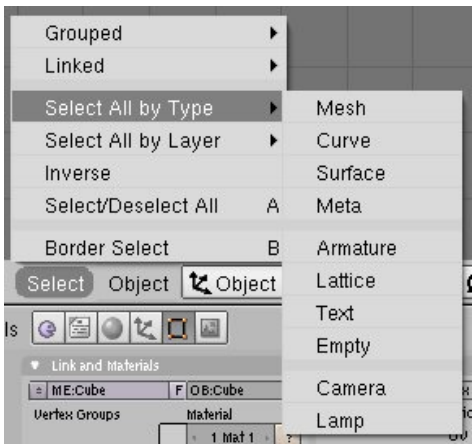
Select All by Type

Mode: Object mode

Hotkey: None

Menu: Select â Select All by Type

Description



Options for selection of objects of one type.
The types are Mesh, Curve, Surface, Meta, Armature, Lattice, Text, Empty, Camera, Lamp.

With this tool it becomes possible to select every **visible** object of a certain type in one go.

Options

Select All by Type in Object mode offers an option for every type of object that can be described by the ObData datablock.

Just take your pick.

Select All by Layer

Mode: Object mode

Hotkey: None

Menu: Select â Select All by Layer

Description

Layers are another means to regroup your objects to suit your purpose.

This option allows the selection of every single object that belongs to a given layer, visible or not, in one single command. This selection is added to anything that was already selected at that moment.



Choice of one layer.

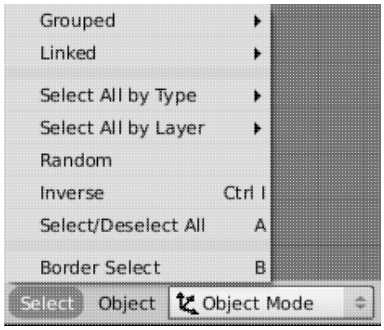
Options

We have the option of selecting the objects of one single layer at a time by LMB -clicking on its number. This has to be repeated for each new layer.

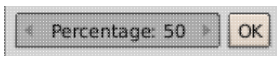
Selection of Objects:

Rather than using the Select All by Layer option, it might be more efficient to make the needed layers visible and use A on them. This method also allows objects to be deselected.

Other Menu Options



Other selecting options.
Available options on the first level of the menu are:



Random Select Percentage.

Random

Randomly selects unselected objects based on percentage probability on currently active layers. On selecting the command a numerical selection box is displayed for the user to select the percentage chance that an object will be selected.
Itâ← s important to note that the percentage represents the likelihood of an unselected object being selected and not the percentage amount of objects that will be selected.

Inverse (CtrlI)
Selects all objects that were not selected while deselecting all those which were.

Select/Deselect All (A)
If anything was selected it is first deselected. Otherwise it toggles between selecting and deselecting every visible object.

Border Select (B)
As described above in the section on border select.

Moving (translating) objects

There are two ways to move or translate an object: moving it by itself, or moving it relative to something else.

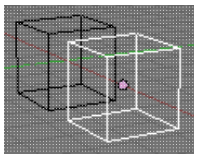
Moving Object(s) Individually

Mode: Object mode

Hotkey: G or Gesture

Menu: Object â Transform â Grab/Move (or Grab/Move on Axis for constraints)

Description



Grab mode.
To translate an object is to place an object in Grab mode. The selected objects will be displayed as white wireframes and can be moved with the mouse (without pressing any mouse buttons); see (*Grab mode*) or keyboard arrow keys.

To confirm the new position, click LMB or press â µ Enter; to cancel Grab mode, click RMB or press Esc. The header of the 3D Window displays the distance you are moving.

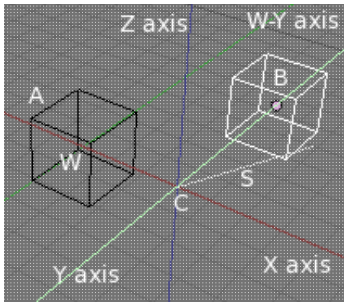
Options

If you keep Ctrl pressed while moving the object you will activate Snap mode, and the object will move by steps (grid squares). Snap mode ends when you release Ctrl so be sure to confirm the position before releasing it. For finer snapping (tenth a grid square) you can hold both Ctrl and â § Shift.

You can control positioning to a finer degree by holding â § Shift while you move. Large mouse movements will translate into very small object movements (tenth the normal amplitude), which allows for finer positioning.

The location of selected objects can be reset to the default value by pressing AltG. The default location is the origin of the global coordinate system.

Axis constraint



Global axes.

Note
For this section, and it sub sections, please reference (*Global axes*).

Movement can be constrained to an axis that is aligned with one of the axes of the global coordinate system, centered on the objectâ← s original world location. The cube â← Bâ← â← s original world location is labeled â← Câ← . The center of the global coordinate system is labeled â← Wâ← ; the Z axis is not visible.

By constraining movement to a global axis you are, in effect, restricting movement to one dimension.

The global aligned axes are color coded as follows:

- X axis is dark red, labeled â← `X axis`â← .
- Y axis is dark green, labeled â← `W-Y axis`â← (for â← world-Y axisâ←).
- Z axis is dark blue, labeled â← `Z axis`â← .

The restricted axis is always highlighted in a lighter shade of color.

For example, the Y axis is drawn in light green if movement is constrained to the Y axis; labeled â← `Y axis`â← .

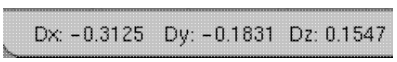
There are two ways to constrain movement: using the mouse or using the keyboard.

Using the mouse

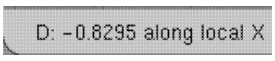
To lock or constrain movement using the mouse, enter Grab mode and move the object while pressing MMB . While in Grab mode you can use the gesture system to pre-select an axis by moving the mouse in a direction roughly inline with a world axis and then clicking and releasing MMB . For example, if you move the mouse along what visually appears to be the X axis and then click and release MMB the objectâ← s movement will be restricted to the world X axis.

Alternately, you can interactively choose the constraining axis by dragging with the MMB while in Grab mode. All three axes become visible with a guide line that emanates from the objectâ← s original location; labeled â← Câ← . This guide is drawn in white dotted line labeled â← Sâ← . As the guide line nears an axis that axis becomes highlighted in a lighter shade and the object snaps to that axis. In this example the guide line is near the Y axis and the cube, labeled â← Bâ← , snaps to it.

Using the keyboard



No Constraint.

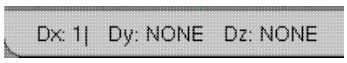


Local Constraint.



Global Constraint.

You can constrain movement to a given axis by pressing either X, Y or Z. A single key press constrains movement to the corresponding *global* axis (*Global Constraint*), as MMB does. A second keypress of the *same* key constrains movement to the corresponding object *local* axis (*Local Constraint*) and a third keypress of the same key removes constraints, (*No Constraint*). The constrained axis is drawn in a lighter color to better visualize the constraint. (*Local Constraint*) and (*Global Constraint*) are all examples of constraints on the X axis using the X key.



Manual entry.
Once grabbing is activated you can enter the Object translation manually by simply typing in a number. This will change the 3D window header as shown in (*Manual entry*).

The number entered is a distance number (i.e. how far from the objectâ← s current location). Think of the â← Dâ← as in â← Displacementâ← , â← Deltaâ← or â← Distanceâ← . The number entered is *not* a world coordinate. To change the objectâ← s world coordinates, use the Transform Properties panel.

By default the X component field is where entry initially goes; see field labeled â← Dxâ← in (*Manual entry*). You can change the default by using the â Tab prior to entering any numbers. For example, to translate 4.4 units along the Y axis you would:

- Enter Grab mode.
- â Tab once.
- Type 4.4.

To translate 3.14 units on the Z axis you would use the â Tab key *twice* prior to entering the number.

Currently you canâ← t delete an incorrect number. You must restart by returning to the original numbers. The â Backspace key will reset to the original values. Hit â µ Enter or Space to finalize and Esc to exit. If you want more flexibility with manual entry use the Transform Properties panel.

It is also possible to enter a value followed by an axis letter to indicate that the value that is entered should be made along the specified axis letter. For example if you wanted to move an object along the Y axis by 3 Blender units you would type GY3â µ Enter or G3Yâ µ Enter. You can also enter negative values to move in the opposite direction.

3 Axis Coordinate/Displacement Entry using â Tab:

In entering X/Y/Z axis numbers at the keyboard, you can use the â Tab to cycle through the X/Y/Z fields that will be altered when a number is entered from the keyboard. It is important to realize that in cycling through the fields you can fill in each field with a different value as you cycle through them, you do not just have to fill in only one field. For example if you wanted to move an object by X:2, Y:3, Z:4, you would type G2 â Tab3 â Tab4 â µ Enter.

As well as being able to constrain along a single specified axis, it is also possible to prevent axis translation/scaling along one axis, but allow translation/scaling along the other two axis. This is achieved by pressing either â § ShiftX, â § ShiftY or â § ShiftZ to prevent translation/scaling along the specified axis.
So if you wished to scale an object on the X and Z axis, but not the Y axis, you could type S â § ShiftY.

Hints

You can use the keyboardâ← s . and the numeric keypadâ← s . NumPad for decimals entry. Be aware that older versions of Blender may not allow the use of the numeric keypadâ← s . NumPad for entering decimals.

Moving/Translating Object(s) by Changing Attributes

Mode: Object mode

Hotkey: CtrlC

Menu: Object â Copy Attributes â (select a set)

Description

Blender has a general purpose way of copying any active objectâ← s attributes to any number of other selected objects. If you copy one objectâ← s location attribute to another object, that second object will be â← movedâ← .

Options

The attributes that can be copied include:

- Location
- Rotation
- Size
- Drawtype
- Time Offset
- Dupli
- Mass
- Damping
- Properties
- Logic Bricks
- Protected Transform
- Object Constraints
- NLA Strips
- Texture Space
- Subsurf Settings
- Modifiers
- Object Pass Index

So, if you â § ShiftRMB -select a cube and a cone, and then â § ShiftRMB -select a lamp last (the lamp thus being the *active* object), and choose CtrlC 1, the cube and the cone will be â← movedâ← to the same location as the lamp.

Rotating objects

Mode: Object mode

Hotkey: R or Gestures

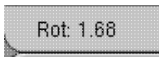
Menu: Object â Transform â Rotate / Rotate on Axis

There are two ways of changing an objectâ← s rotation; individually, and by copying the rotation attribute from another object as described above.

Description

Change the rotation by moving the mouse and confirming with LMB or â µ Enter. You can cancel with RMB or Esc.

Rotation in 3D space occurs around an axis, and there are several ways to define this axis. But in general an axis is defined by a direction line and a point that the line passes through. By default the axis is orthogonal to your screen (i.e. it is going into or out of your screen). If you are viewing the scene from the front, side, or top 3D view windows, the rotation axis will be parallel to one of the global coordinate system axes. If you are viewing the scene from an angle, the rotation axis is angled too, which can easily lead to a very odd rotation of your object. In this case, you may want to keep the rotation axis parallel to the coordinate system axes.



Rotation in header.
As you rotate the object the angle of rotation is displayed in the 3D window header (*Rotation in header*).

Options

As with translation, for finer control or precision, use Ctrl and/or â § Shift. Pressing Ctrl switches to Snap mode, and rotations are constrained to 5 degree increments. Pressing â § Shift at the same time constraints the rotation to 1 degree increments. Pressing â § Shift alone while rotating allows finer degrees of rotation as precise as 1/100th of a degree.

The rotation of selected objects can be reset to the default value by pressing AltR.

Axis of rotation Constraint

Just like Grab mode you can constrain the axis of rotation by using either the mouse or the keyboard. The only difference is that you only enter an angle. See Grab modeâ← s axis constraint description for exact details.

Point of rotation

To select the point of rotation (so-called â← pivot pointâ←) that the rotation axis will pass through, use the â← rotation/scalingâ← button accessed in the header of the 3D window, . For a full description of the different pivot points, see their dedicated page.

Hints

If youâ← re just getting started with rotation, donâ← t worry too much about the foregoing details. Just play around with the tool and youâ← ll get a feeling for how pivot points affect rotation. For example, an easy way to understand how pivot points work is to create two cubes, and then cycle through each pivot point type while in Rotate mode.

For example, to have one cube orbit another cube select the Active Object pivot point. As you rotate, constrained or not, the other object(s) orbit the active object.

Scaling objects

Mode: Object mode

Hotkey: S or Gesture

Menu: Mesh â Transform â Scale

Description

Scale the objects by moving the mouse and confirming with LMB or â µ Enter, and cancel with RMB or Esc.

Scaling in 3D space occurs around a center point; much like a rotation occurs around a pivot point. If you increase the size of the object, all points are moved away from the selected center point; if you decrease it, all points move towards this point.

Options

Here again the Ctrl key switches to Snap mode, with discrete scaling at **0.1** steps. Press â § Shift for fine tuning.

The scaling of selected objects can be reset to the default value by pressing AltS.

Axis of scale Constraint

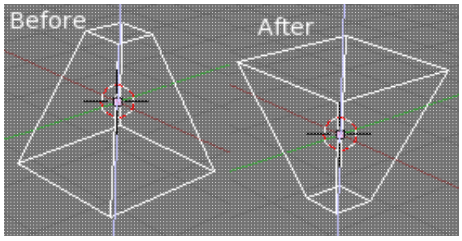
By default, the selected objects are uniformly scaled in all directions. To change the proportions (make the object longer, broader and so on), you can lock the scaling process to one of the global coordinate axes, just as you would with Grab mode and Rotate mode. Again all considerations on constraining to a specific axis, in respect to Grabbing, still hold as well as those on numerical input. See Grab modeâ← s axis constraint description for exact details.

Center point of scale

To select the center point of scale (so-called â← pivot pointâ← !) use the â← rotation/scalingâ← button accessed in the header of the 3D window, . For a full description of the different pivot points, see their dedicated page.

Mirroring objects

Mirroring objects is a different application of the scale tool. Mirroring is effectively nothing but scaling with a negative factor in one direction. For example, to mirror in the direction of any single axis:



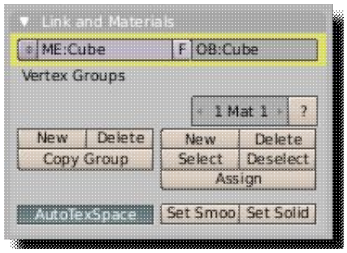
Mirrored Frustum.

- Enter Scale mode.
- Select an axis using X, Y or Z key.
- Enter â← -1â← as the scaling factor.

(*Mirrored Frustum*) is an example of mirroring a frustum object along the Z axis. These are the steps to mirror the frustum:

- Enter Scale mode.
- Select the Z axis using the Z key.
- Enter â← -1â← as the scaling factor.
- Hit â µ Enter.

Skinning and Cloning Objects



Relevant fields highlighted in yellow.
At the very top of the Link and Materials panel, you will find two fields, one in light pink and another right next to it in gray.

The field in gray starts with OB: and is the name of the object itself. It has to be unique within the .blend file across all scenes. The field name on the left starts with a two-letter abbreviation indicating what type of object it is, and the name of its skin, or physical appearance:

- ME: is the physical mesh, made up of vertices.
- CU: is a curve, surface, or text object, made up of control points.
- MB: is a metaball, whose skin is represented as a mathematical function.

Any of these skins can be shared by objects. Imagine a scene with 50 cats, some skinny, some fat. You would have two meshes, ME:`Cat.Skinny` and ME:`Cat.Fat`. You would create 50 OB:`Cat.001`, OB:`Cat.002`, â←¦, OB:`Cat.050`, and assign 20 of the OBjects to be fat cats, and the rest skinny.

Options

Clicking the F will fake a user of the skin, and it will not be deleted when no one uses it. The next time you open the .blend file, it will be in memory and will not have to be re-made. You can then create an object of its type, and use that skin.

At any time you can change the skin of an object by clicking the up-down selector on the left of the field and selecting a different skin for that same object type. When you do, the field will then show the multi-user button, â← 2â← identifying how many other objects share this skin.

Hotkey: AltD

Menu: Object â Duplicate Linked

Select an object and use the hotkey to create a clone of the original. The two objects will share the same skin. This means that altering either object at the Edit mode level (when in Edit mode), by for example grabbing vertices, will result in the other objects being altered in the same relative way. This linking usually only works when in Edit mode, so scaling/rotating/grabbing an object in Object mode will not result in the other linked object being affected.

For more details on datablocks, duplication and linking, and so on, see the Data System chapter.

Regrouping Objects

There can be many objects in a scene: A typical stage scene consists of furniture, props, lights, and backdrops. Blender helps you keep everything organized by allowing you to group like objects together.

When modeling a complex object, such as a watch, you may choose to model the different parts as separate objects. However, all of the parts may be attached to each other. In these cases, you want to designate one object as the parent of all the children. Movement and rotation of the parent also affects the children.

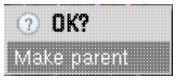
Parenting objects

Mode: Object mode

Hotkey: CtrlP

Menu: Object â Parent â Make Parent

Description



Make Parent.
To parent objects, select at least two objects (Select the child objects first, and select the parent object last), and press CtrlP. A confirmation dialog will pop up asking Make Parent. Selecting Make Parent confirms and the child/children to parent relationship is created, see image (*Make Parent*). The last object selected will be the *active* object (outlined in pink), and will also be the *parent* object. If you selected multiple objects before selecting the parent, they will all be children of the parent and will be at the same level of the hierarchy (they are â← siblingsâ←).

Moving and rotating the parent will also usually move/rotate the child/children. However moving/rotating the child/children of the parent, will not result in the parent moving/rotating. In other words, influence is usually descendant (parent â child/children), and not ascendant (child/children â parent).

Mode: Edit mode

Hotkey: CtrlP

Menu: Mesh â Vertices â Make Vertex Parent

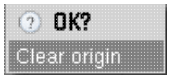
You can parent an object to a single vertex or a group of vertices as well; that way the child/children will move when the parent mesh is deformed, like a mosquito on a pulsing artery. In Object mode, select the child/children and then the parent object. â Tab into Edit mode and on the parent object select either one vertex that defines a single point, or select three vertices that define an area (the three vertices do not have to form a complete face they can be any three vertices of the parent object), and then hit CtrlP and confirm.

At this point if a single vertex was selected a relationship/parenting line will be drawn from the vertex to the child/children. If three vertices were selected then a relationship/parenting line is drawn from the averaged center of the three points (of the parent object) to the child/children. Now, as the parent mesh deforms and the chosen parent vertex/vertices move, the child/children will move as well.

Note
It is in fact a sort of â← reversedâ← hookâ←¦

Options

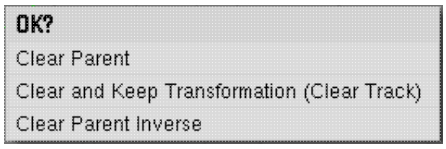
Move child



Move child.
You can *move* a child object to its parent by clearing its origin. The relationship between the parent and child isnâ← t affected. Select the child object and press AltO. By confirming the dialog the child object will snap to the parentâ← s location. Use the Outliner view to verify that the child object is still parented.

Remove relationship/Clear Parent

You can *remove* a parent-child relationship via AltP; see image (*Remove relationship*).



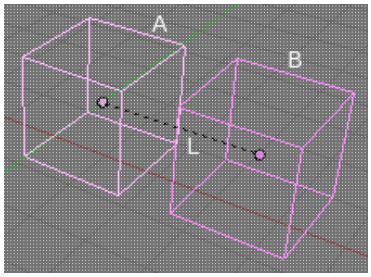
Remove relationship.
The menu contains:

Clear Parent
: If the parent in the group is selected nothing is done. If a child or children are selected they are disassociated with the parent, or freed, and they return to their *original* location, rotation, and size.

Clear and Keep Transformation (Clear Track)
: Frees the children from the parent, and *keeps* the location, rotation, and size given to them by the parent.

Clear Parent Inverse
: Places the children with respect to the parent as if they were placed in the Global reference. This effectively clears the parentâ← s transformation from the children. For example, if the parent is moved 10 units along the X axis and Clear Parent Inverse is invoked, any selected children are freed and moved -10 units back along the X axis. The â← Inverseâ← only uses the last transformation; if the parent moved twice, 10 units each time for a total of 20 units, then the â← Inverseâ← will only move the child back 10 units, not 20.

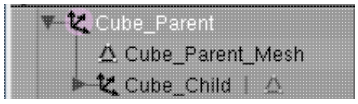
Parenting Example



Parenting Example.
The active object, in light pink (cube A), will be made the parent of all the other object(s) in the group (darker pink/purple cube B). The center(s) of all children object(s) are now linked to the center of the parent by a dashed line; see image (*Parenting Example*). The parent object is cube â← Aâ← and the child object is cube â← Bâ← . The link is labeled â← Lâ← .

At this point, grabbing, rotating, and scaling transformations to the parent will do the same to the children. Parenting is a very important tool with many advanced applications, as weâ← ll see in later chapters; it is used extensively with advanced animations.

Hints



Outliner view.
There is another way to see the parent-child relationship in groups and that is to use the Outliner view of the Outliner window. Image (*Outliner view*) is an example of what the Outliner view looks like for the (*Parenting Example*). Cube â← Aâ← â← s object name is â← `Cube_Parent`â← and cube â← Bâ← is â← `Cube_Child`â← .

Separating Objects

Mode: Edit mode

Hotkey: P

Menu: Mesh â Vertices â Separate

Description

At some point, youâ← ll come to a time when you need to cut parts away from a mesh to be separate, but you might wonder how to do that. Well, the operation is easy.

To separate an object, the vertices (or faces) must be selected and then separated, though there are several different ways to do this.

Options



Suzanne dissected neatly.

Selected
: This option separates the selection to a new object.

All Loose Parts
: Separates the mesh in its unconnected parts.

By Material
: Creates separate mesh objects for each material.

Grouping objects

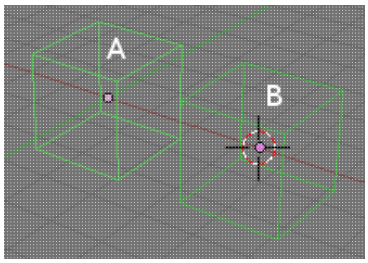
Mode: Object mode

Panel: Object â Object and Links

Hotkey: CtrlG

Menu: Object â Parent â Add to New Group

Description



Grouped objects.
Group objects together without any kind of transformation relationship. Use groups to just logically organize your scene, or to facilitate one-step appending or linking between files or across scenes. Objects that are part of a group always shows as light green when selected; see image (*Grouped objects*).

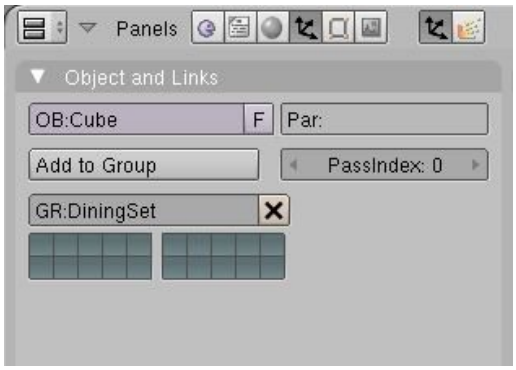
Options



Groups pop-up menu.

Adding to, or Creating, a Group

CtrlG pops up a dialog for adding to existing groups or creating a new group; see image (*Groups pop-up menu*). This same menu is also available *via* the 3D view header: Object â Group.
Alternatively, with the object selected or in Edit mode, click the Add to Group button shown below in image (*Naming a Group*). The popup list allows you to click on an existing group, or create a new one.
You can also ostracize, or banish, the selected object from all groups by selecting the Remove from All Groups option.



Naming a Group.

Naming a Group

To name a group, you can use the Buttons window, Object context (F7), Object and Links panel: just click â § ShiftLMB in the GR: field and type a meaningful name.
To name groups in the Outliner window, select Groups as the outliner display from the header combo box, and CtrlLMB click on the group name. The name will change to an editable field; make your changes and press â µ Enter.

Restricting Group Contents via Layers

That cluster of layer buttons below a group designation determines from which layers the group objects will be included when duplicated. If your group contains objects on layers 10, 11 and 12, but you disable the layer 12 button in the group controls, duplicates of that group (using the Dupligroup feature) will only show the portions of the group that reside in layers 10 and 11.

Appending or Linking Groups

To append a group from another .blend file, consult this page. In summary, File â Append or Link â `filename` â Group â `groupname`.

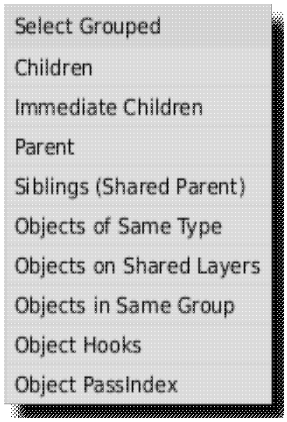
Select Grouped

Mode: Object mode

Hotkey: â § ShiftG

Menu: Select â Grouped

Description



Selected Grouped pop-up menu.
â § ShiftG pops up a dialog for selecting objects based on parenting and grouping characteristics; see image (*Selected Grouped pop-up menu*).

Options

Children
: Selects all the active objectâ← s children, and the childrenâ← s children, up to the last generation.

Immediate Children
: Selects all the active objectâ← s children, but not those of these children (no â← grandchildrenâ← selected!).

Parent
: Selects the parent of the active object and deselects the active object.

Siblings (Shared Parent)
: Selects all the siblings of the active object.

Objects of Same Type
: Selects objects based on the current active object type.

Objects on Shared Layers
: Selects all objects on the same layer(s) of the active object.

Objects in Same Group
: Selects objects that belong to the same group as the active object.

Object Hooks

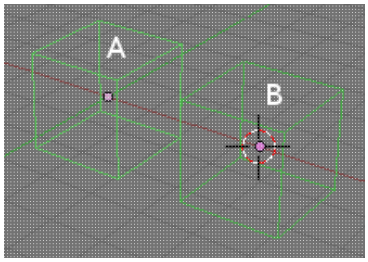
Selects all hooks which are attached to the active object.

Object PassIndex
: Selects all objects which have the same PassIndex number as the active object. See the ID Mask node usage for more information on this option.

Object Color
: Selects objects that have the same color (the one set in Draw panel) as the active object.

Game Properties
: Selects objects that have the same game properties as the active object.

Example



Grouped objects.
(*Grouped objects*) shows two cubes grouped together where A is the last selected object indicated by being drawn in a lighter color.

Tracking

Mode: Object mode

Panel: Object â Constraints

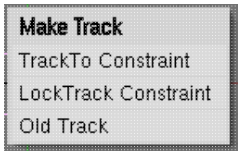
Hotkey: CtrlT

Menu: Object â Track â Make Track

Description

Tracking consists of one object watching another. The watcher is the â← Trackerâ← and the watched is the â← Targetâ← . If the target moves the tracker rotates; if the tracker moves the tracker rotates. In both cases the tracker maintains a constant heading towards the target.

Options



Make Track menu.
To make one or more objects track another object (the target), select at least two objects and press CtrlT. The active object becomes the target and the others objects the trackers. The (*Make Track menu*) provides several options for creating the initial tracking:

TrackTo Constraint

The Track To constraint applies rotations to the object, so that it always points a local â← Toâ← axis towards the target object, with another local â← Upâ← axis permanently maintained as vertical as possible. This tracking is similar to the billboard tracking in 3D. This is the preferred tracking constraint, because it has a more easily controlled constraining mechanism.

The â← billboard trackingâ←
The term â← billboardâ← has a specific meaning in real-time CG programming (i.e. video games!), where it is used for plane objects always facing the camera (they are indeed â← trackersâ← , the camera being their â← targetâ←). Their main usage is as support for tree or mist textures: if they werenâ← t permanently facing the camera, you would often see your trees squeezing to nothing, or your mist turning into a millefeuille paste, which would be funny but not so credibleâ←¦

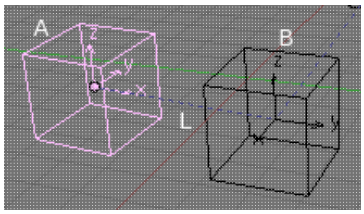
This constraint shares a close relationship to the IK constraint in some ways. This constraint is very important in rig design, and you should be sure to read and understand the tracking tutorial, as it centers around the use of both of these constraints.

The controls for changing the â← Trackingâ← and â← Upâ← axis of a tracker object are located in the Constraints panel, in Buttons window, Object context (F7), Object sub-context. The constraint fields are:

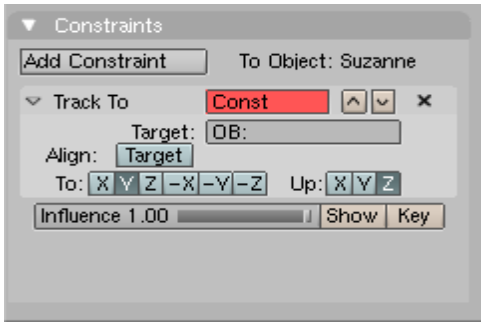
To
: The axis of the object that has to point to the target.

Up
: The axis of the object that has to be aligned (as much as possible) with the world Z axis. An Align: Target button, when enabled, uses the coordinates of the target object's Z axis, thus tilting or rocking the object as it tracks the target.

Head/Tail - With Bone targets only
: A number from 0.0 to 1.0 that represents the place on the target bone to Track to (0.0 = the bone's root; 1.0 = the bone's head)



TrackTo example.



The Constraints panel of the example.

(*TrackTo example*) is an example of a cube using the TrackTo constraint. Cube â← Aâ← is tracking â← Bâ← where â← Lâ← is the tracking line. Notice how the tracking objectâ← s local axes are visible by using the Draw panelâ← s axis button. You can clearly see the tracking â← Toâ← and â← Upâ← axis. Cube â← Aâ← â← s constraint setting are reflected in (*The Constraints panel of the example*). +X is the â← Toâ← axis and Z is the â← Upâ← axis.

You can also see in (*The Constraints panel of the example*) what cube â← Aâ← is tracking by looking at the Target field. We can see that cube â← Aâ← is tracking cube â← Bâ← because cube â← Bâ← â← s name is â← Cubeâ← . You can redirect tracking to another object simply by entering in the name of another object.

LockTrack Constraint

Locked Track is a difficult constraint to explain, both graphically and verbally. The best real-world example would have to be a compass. A compass can rotate to point in the general direction of its target, but it canâ← t point *directly at* the target, because it spins like a wheel on an axle. If a compass is sitting on a table and there is a magnet directly above it, the compass canâ← t point to it. If we move the magnet more to one side of the compass, it still canâ← t point *at* the target, but it can point in the general direction of the target, and still obey its restrictions of the axle.

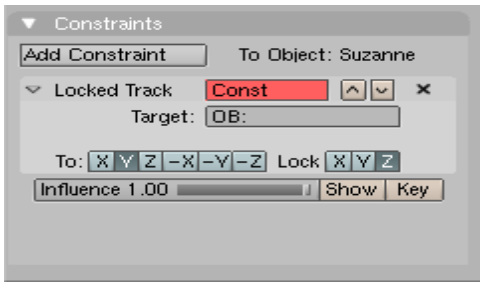
When using a Locked Track constraint, you can think of the target object as a magnet, and the affected object as a compass. The â← Lockâ← axis will function as the axle about which the object spins, and the â← Toâ← axis will function as the compass needle. Which axis does what is up to you! If you have trouble understanding the buttons of this constraint, read the tool-tips; theyâ← re pretty good. If you donâ← t know where your objectâ← s axes are, turn on the Axis button in the Draw panel, Object context (F7). Or, if youâ← re working with bones, turn on the Draw Axes button, Armature panel, Editing context (F9).

This constraint was designed to work cooperatively with the Track To constraint. If you set the axes buttons right for these two constraints, Track To can be used to point the axle at a target object, and Locked Track can spin the object around that axle to a secondary target.

This constraints also works very well for 2D billboarding.

This is all related to the topic discussed at length in the tracking tutorial.

The controls for changing the tracking â← Toâ← and â← Lockâ← axis are located in the Constraints panel. This panel is located in the same place as the Anim settings and Draw panels (Object context, F7). The constraint fields are:



The Locked Track Constraints panel.

Target
: The name of the target object that the tracking object tracks.
VG
: This is only available when tracking to a mesh object. If you type in the name of a vertex group of the target, it will be the center of this vertex group that will be the â← physicalâ← target, not the objectâ← s center.
BO
: This is only available when tracking to an armature object. If you type in the name of a bone of the target, it will be this bone that will be the â← physicalâ← target, not the armatureâ← s center.

To
: The tracking axis. It shouldnâ← t be the same as the Lock axis.

Lock
: The locked local axis.

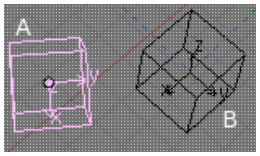
Influence
: This controls how accurately the tracking object tracks the target. **0** means that the constraint is turned off. The tracking object will remain locked in orientation. **1** means tracking is completely on and the tracking axis will stay tightly focused on the target.

Show
: This adds an influence Ipo channel to the constraint if one is not present. You can then add keys to the channel.

Key
: This adds animation keys to the influence Ipo channel. This is a very powerful combination. For example, you could have a camera with a Locked Track constraint applied and have input driving the influence channel.

For further details see TrackTo constraint. LockTrack and TrackTo are very similar where the former has a â← Lockâ← ed axis verses an â← Upâ← axis.

Old Track



Old Track â← constraintâ← .
This is an older algorithm prior to version 2.30, and is similar to Track To constraint in that no axis is locked. This algorithm merely tries to keep a â← Toâ← axis pointed at the target. The tracking object will usually end up in an odd orientation when this constraint is first applied. In order to get correct results use AltR when applying or changing the tracking or â← Upâ← axis. However, the preferred method to use is Track To constraint.

Letâ← s assume you have nonetheless selected Old Track in the dialog with two cubes selected; see (*Old Track â← constraintâ←*). By default the inactive object(s) track the active object so that their local +Y axis points to the tracked object. Cube â← Aâ← is tracking cube â← Bâ← using the Old Track constraint. You can see that â← Aâ← â← s +Y axis is pointing at â← Bâ← but at an odd orientation. This typically happens if the object already has a rotation of its own. You can produce correct tracking by canceling the rotation on the tracking object using AltR.

The orientation of the tracking object is also set such that the chosen â← Upâ← axis is pointing upward.



Setting track axis.
If you want to change this you need to get to the Anim settings panel where Old Trackâ← s settings are accessed. First select the tracking object (not the target) and change the Button window to Object context by clicking the icon (), or F7; see (*Setting track axis*).

You then have the option of selecting the *Tracking axis* from the first column-set of six radio buttons and/or selecting the *upward-pointing* axis from the second column-set in the Anim Setting panel. Each time you change the â← Upâ← axis you need to apply AltR otherwise the tracking object will continue to track with the old orientation. This is one of the drawbacks to using Old Track.

To clear or remove an old track â← constraintâ← , select the tracking object and press AltT. As with clearing a parent constraint, you must choose whether to lose or save the rotation imposed by the tracking.

Note
AltT command works (and is useful) only for the Old Track â← constraintâ← . To clear the Track To and Locked Track constraints, just delete them directly from the stack at the Constraints panel.

Hints

The *active* object always becomes the target object to be tracked. In all but Old Track a blue dashed line is drawn between the tracker and target indicating that a tracking constraint is in place between the

corresponding objects. If you see an object tracking another object without a dashed blue line then you know the tracking object is using the Old Track â← constraintâ← .

Invalid Tracking or settings

If you choose an invalid tracking â← Toâ← and/or â← Upâ← axis, the tracking object keeps it current orientation and ignores the incorrect selections. For example, if you choose the +Z axis as the â← Toâ← axis and also choose the +Z axis as the â← Upâ← axis, you have chosen an invalid combination because you canâ← t have the tracking objectâ← s +Z axis doing two different things at the same time.

If you have problems setting up the correct â← Toâ← and â← Upâ← axes you may want to turn on the tracking objectâ← s local axes. You can do this from the Draw panel by clicking on the Axis button. See The Interface chapter for further details on the Draw panel.

Duplicate

Mode: Edit and Object modes

Hotkey: â § ShiftD

Menu: Object â Duplicate

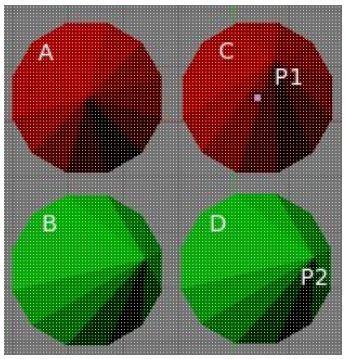
Description

This will create a visually identical copy of the selected object(s). The copy is created at the same position as the original object and you are automatically placed in Grab mode. Reference (*Duplicate Example*) for the discussions below.

This copy is a new object, which â← **shares**â← some datablocks with the original object (by default, all the Materials, Textures, and Ipos), but which has copied others, like the mesh for example. This is why this form of duplication is sometimes called â← shallow linkâ← , because all datablocks arenâ← t shared, some of them are â← hard copiedâ← !

Note that you can choose which types of datablock will be linked or copied when duplicating: in the User Preferences window (should be the one at the top of your screen), click on the Edit Methods â← tabâ← , and in the Duplicate with object: buttons group, activate those corresponding to the types of datablocks you want to really copy - the others will just be linked.

Examples



Duplicate Example.
The cone labeled â← Câ← is a *Duplicate* of cone â← Aâ← . Here are some properties to notice:

- The vertex at â← P1â← has been moved but the same vertex on cone â← Aâ← is unchanged. This means the mesh data are copies, not links.
- Cone â← Câ← â← s color is red because cone â← Aâ← â← s color is red. This means the material properties are linked not copied.
- If you rotate cone â← Câ← , cone â← Aâ← remains unchanged. This means the transform properties (i.e. Object datablocks) are copies, not links.

See above if you want separate copies of the datablocks normally linked.

Linked Duplicates

Mode: Edit and Object modes

Hotkey: AltD

Menu: Object â Duplicate Linked

Description

You also have the choice of creating a *Linked Duplicate* rather than a *Duplicate*; this is called a deep link. This will create a new object with **all** of its data linked to the original object. If you modify one of the linked objects in Edit mode, all linked copies are modified. Transform properties (object datablocks) still remain copies, not links, so you still can rotate, scale, and move freely without affecting the other copy. Reference (*Duplicate Example*) for the discussions below.

Examples

The cone labeled â← Dâ← is a *Linked Duplicate* of cone â← Bâ← using AltD. Here are some properties to notice:

- The vertex at â← P2â← has moved and the same vertex on cone â← Bâ← has moved as well. This means the mesh data are links, not copies.
- Cone â← Dâ← â← s color is green because cone â← Bâ← â← s color is green. This means the material properties are also linked and not copied.
- If you rotate cone â← Dâ← cone â← Bâ← remains unchanged. This means the transform properties are copies, not links.

A common table has a top and four legs. Model one leg, and then make linked duplicates three times for each of the remaining legs. If you later make a change to the mesh, all the legs will still match. Linked duplicates also apply to a set of drinking glasses, wheels on a carâ←¦ anywhere there is repetition or symmetry.

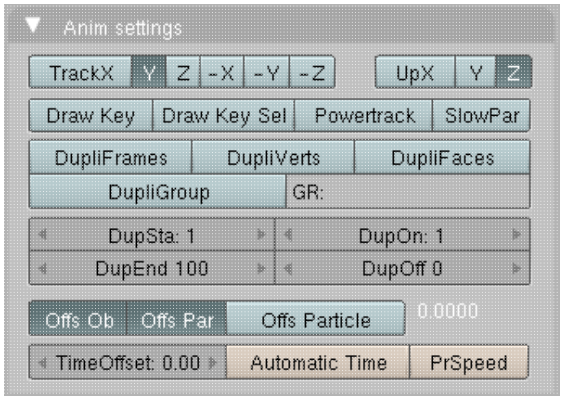
Procedural Duplication

Mode: Object mode and Edit mode

Panel: Anim settings

Hotkey: F7

There are currently four ways in Blender to procedurally duplicate objects. These options are located in the Object context (F7), panel Anim settings.



Anim settings panel.

DupliVerts
: This creates an instance of all children of this object on each vertex (for mesh objects only).

DupliFaces
: This creates an instance of all children of this object on each face (for mesh objects only).

DupliGroup
: This creates an instance of the group with the transformation of the object. Group duplicators can be animated using actions, or can get a Proxy.

DupliFrames
: For animated objects, this creates an instance on every frame. As youâ← ll see on this topicâ← s subpage, this is also a *very* powerful technique for arranging objects and for modeling them.

Hints

If you want transform properties (i.e. object datablocks) to be â← linkedâ← , see the page on parenting.

DupliVerts

Mode: Object mode

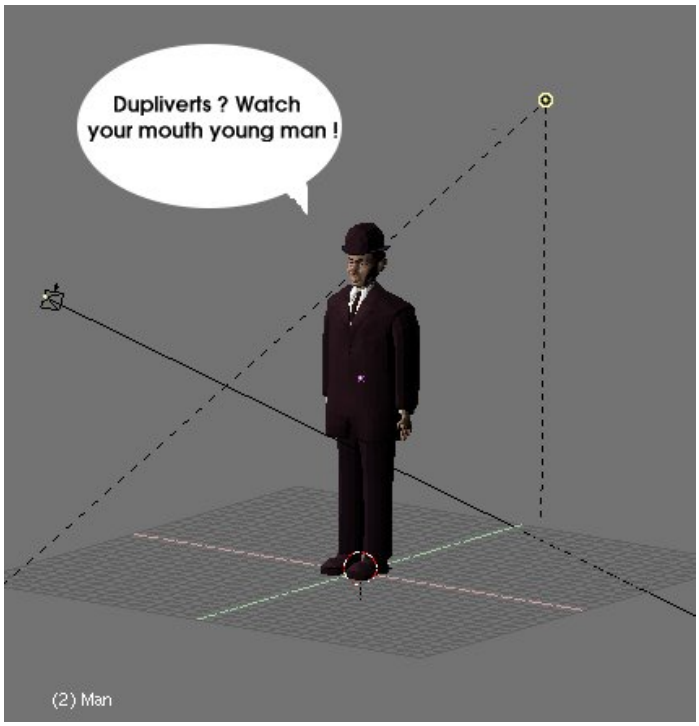
Panel: Anim Settings â DupliVerts

Hotkey: F7

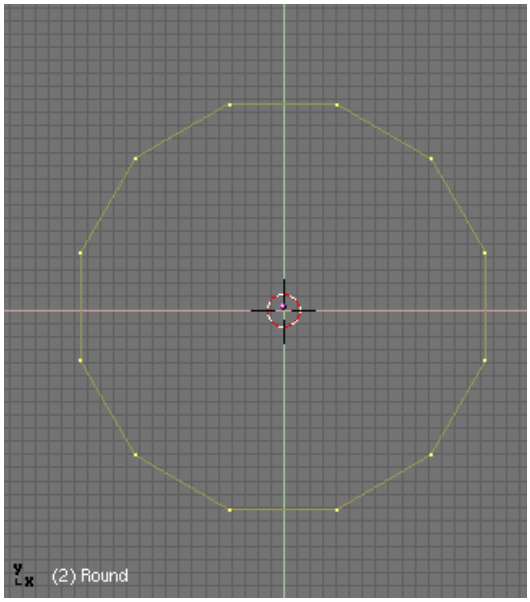
DupliVerts are not a rock band nor a dutch word for something illegal (well maybe it is) but is a contraction for *Dupli*cation at *Vert*ices, meaning the duplication of a base object at the location of the vertices of a mesh (or even a particle system). In other words, when using DupliVerts on a mesh, an instance of the base object is placed on every vertex of the mesh.

There are actually two approaches to modeling using DupliVerts. They can be used as an arranging tool, allowing us to model geometrical arrangements of objects (e.g. the columns of a Greek temple, the trees in a garden, an army of robot soldiers, the desks in a classroom). The object can be of any object type which Blender supports. The second approach is to use them to model an object starting from a single part of it (i.e.: the spikes in a club, the thorns of a sea-urchin, the tiles in a wall, the petals in a flower).

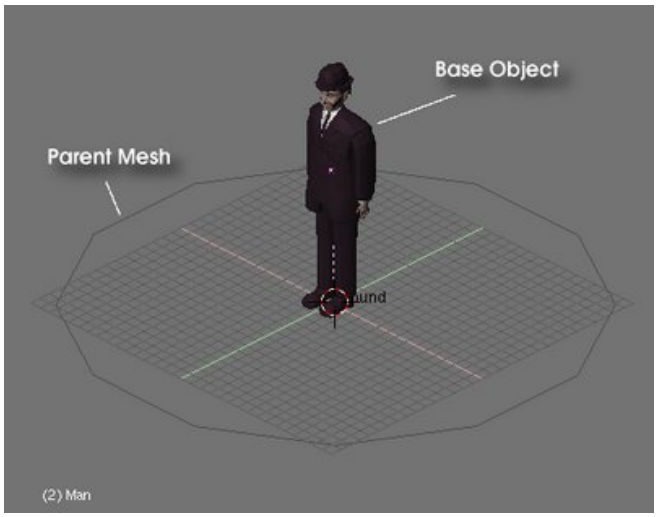
DupliVerts as an Arranging Tool



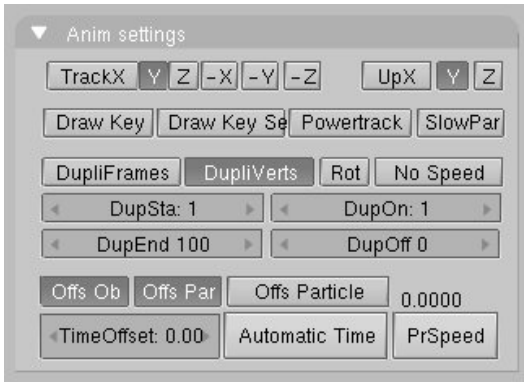
A simple scene to play with.
All you need is a base object (e.g. the *tree* or the *column*) and a pattern mesh with itâ← s vertices following the pattern you have in mind. In this section, we will use a simple scene for the following part. It consists of a camera, the lamps, a plane (for the floor) and a strange man I modeled after Magritteâ← s famous character (*A simple scene to play with*). If you donâ← t like surrealism you will find this part extremely boringâ←¦



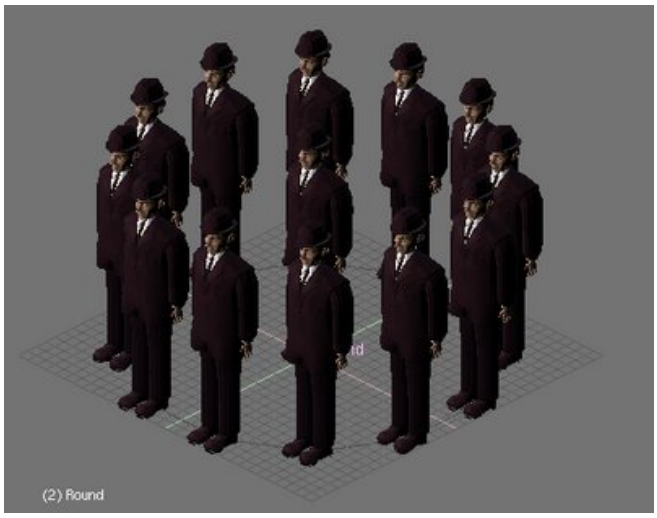
A circle for a parent mesh.
Anyway, the man will be my *base object*. It is a good idea that he will be at the center of the coordinate system, and with all rotations cleared. Move the cursor to the base objectâ← s center (this one being selected, do â § ShiftS â Cursor -> Selection), and from Top view add a mesh (the example uses a circle as a pattern, with 12 vertices or so (*A circle for a parent mesh*). The pattern object can be a two-dimensional primitive (plane or circle), or even a three-dimensional primitive mesh (cube, tube, sphere) or a curve (two dimensional, or a three-dimensional path), or even your own custom mesh, so long as it has vertices (you cannot use a camera, for example, but you can use a large landscape mesh if you want to plant trees - trees being your base object).



The man is parented to the circle.
In Object mode, select the base object and add the circle to the selection (order is very important here). Parent the base object to the circle by pressing CtrlP. Now, the circle is the parent of the character (*The man is parented to the circle*).

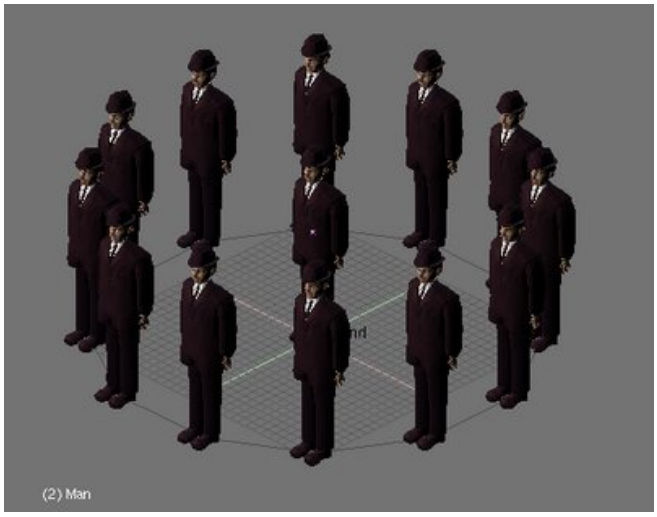


Anim settings panel.
Now select only the circle, switch the Buttons window to the Object context (via or F7) and select the DupliVerts button in the (*Anim settings panel*).

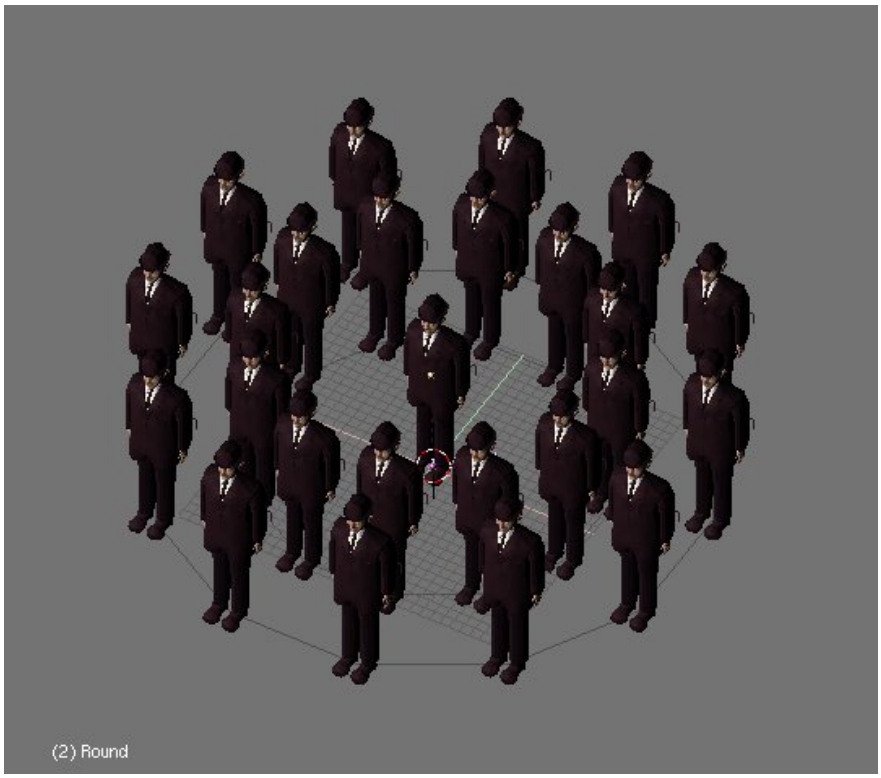


In every vertex of the circle a man is placed.
Wow, isnâ← t it great? Donâ← t worry about the object at the center (*In every vertex of the circle a man is placed*). It is still shown in the 3D views, but it will *not* be rendered. You can now select the base object, edit it in Edit mode (scale, rotate,â←¦â← be careful with the Object mode: scaling, among others, might bring up some problems when applying rotation to DupliVerts, as we will see soon), and all DupliVerted objects will reflect the changes. But the more interesting thing to note is that you can also edit the parent circle.

Note
The base object is not rendered if DupliVerted on a mesh but it *is* rendered if DupliVerted on a particle system! This doesnâ← t appear to be true on Blender 2.45 and later.



Changing the size of the circle in Edit mode.
Select the circle and scale it. You can see that the mysterious men are uniformly scaled with it. Now enter the Edit mode (â Tab) for the circle, select all vertices (A) and scale it up about three times. Leave Edit mode and the DupliVerted objects will update (*Changing the size of the circle in Edit mode.*). This time they will still have their original size but the distance between them will have changed. Not only can we scale in Edit mode, but we can also delete or add vertices to change the arrangement of men.

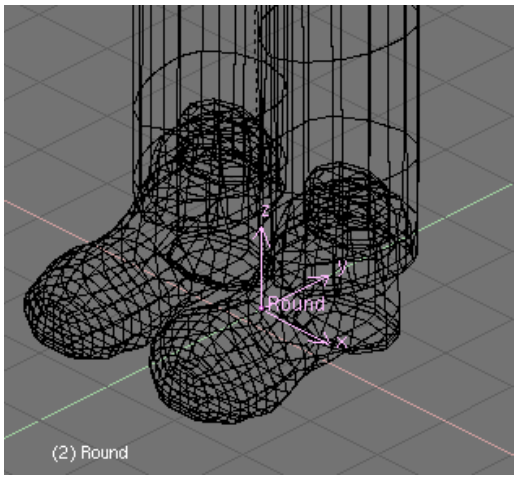


A second row of Magritteâ ← s men.
Select all vertices in Edit mode and duplicate them (â § ShiftD). Now scale the new vertices outwards to get a second circle around the original. Leave Edit mode, and a second circle of men will appear (*A second row of Magritteâ ← s men.*).

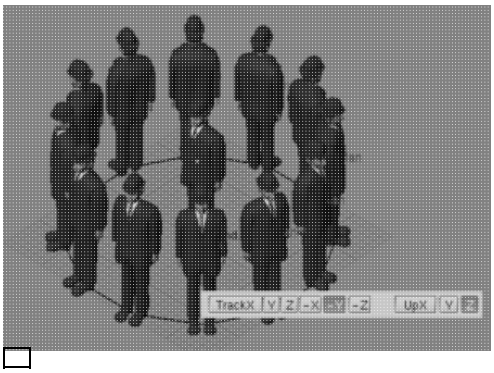
Until now all Magritteâ ← s men were facing the camera, ignoring each other. We can get more interesting results using the Rot button next to the DupliVerts button in the Anim settings panel. With this toggle button active, we can rotate the DupliVerted objects according to the normals of the parent object. More precisely, the DupliVerted objects axis are aligned with the normal at the vertex location. Which axis is aligned (X, Y or

Z) with the parent mesh normal depends on what is indicated in the TrackX, Y, Z buttons and the UpX, Y, Z buttons top in the Anim settings panel. Trying this with our surrealist buddies, will lead to weird results depending on these settings.

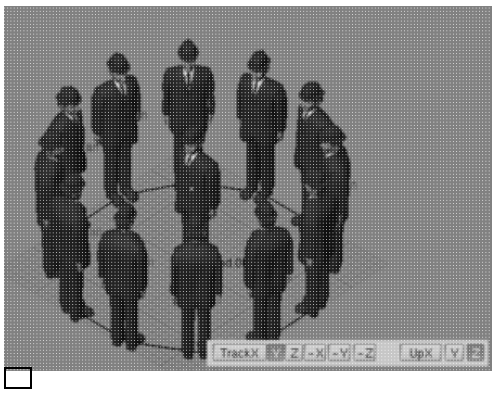
The best way to figure out what will happen is first of all aligning the â← baseâ← and â← parentâ← objectsâ← axis with the World axis. This is done selecting both objects and pressing CtrlA, and click the Apply Size/Rot? menu.



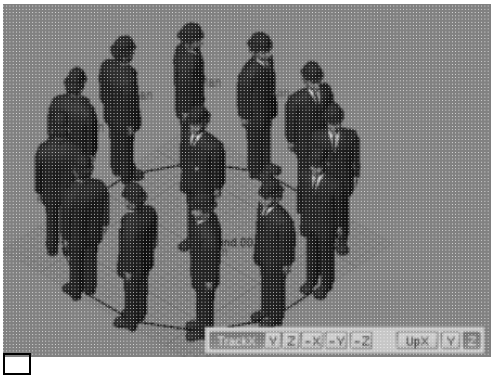
Show objectâ← s axis to get what you want.
Then make the axis of the base object and the axis and normals in the parent object visible (*Show objectâ← s axis to get what you want* â← in this case, being a circle with no faces, a face must be defined first for the normal to be visible â← actually to exist at all). Now select the base object (our Magritteâ← s man) and play a little with the tracking buttons. Note the different alignment of the axis with the different combinations of UpX, Y, Z and TrackX, Y, Z (*Negative Y Axis is aligned to vertex normal (pointing to the circleâ← s center), Positive Y axis is aligned to normal, Positive X axis is aligned to normal, Positive Z axis is aligned to normal (weird, huh?)*).

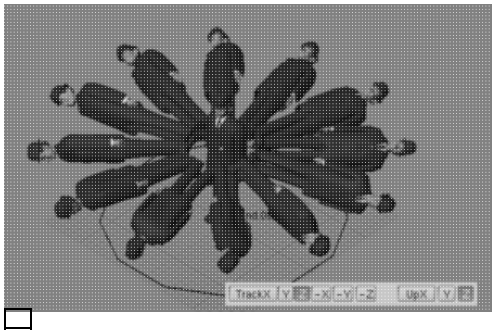


Negative Y Axis is aligned to vertex normal (pointing to the circleâ← s center).



Positive Y axis is aligned to normal.



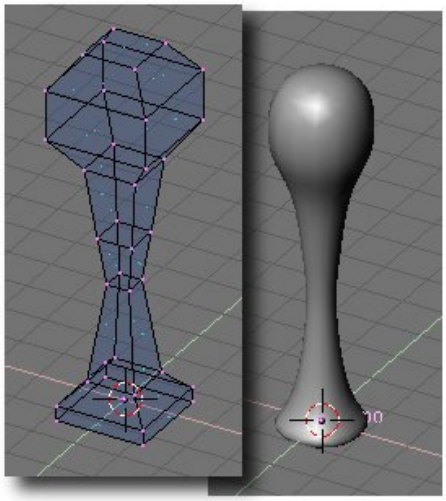


Positive X axis is aligned to normal.

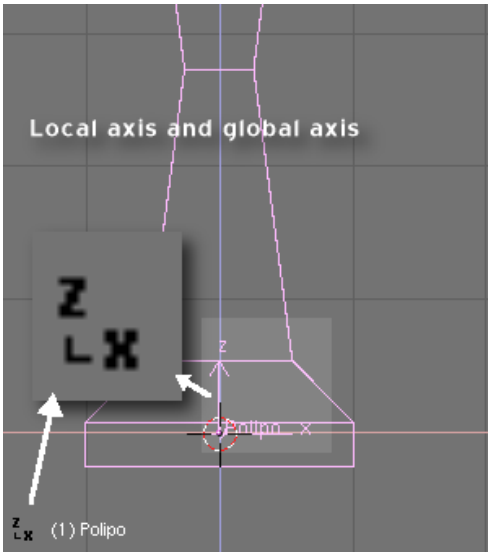
Positive Z axis is aligned to normal (weird, huh?).

DupliVerts to Model a Single Object

Very interesting models can be made using DupliVerts and a standard primitive. Starting from a cube in Front view, and extruding a couple of times I have modeled something which looks like a tentacle when SubSurfs are activated (*Strange tentacle and SubSurfed version*). Then I added an icosphere with 2 subdivisions.



Strange tentacle and SubSurfed version.

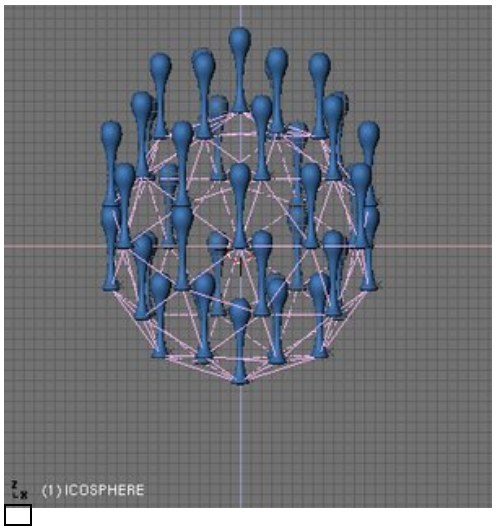


Local reference of the tentacle.

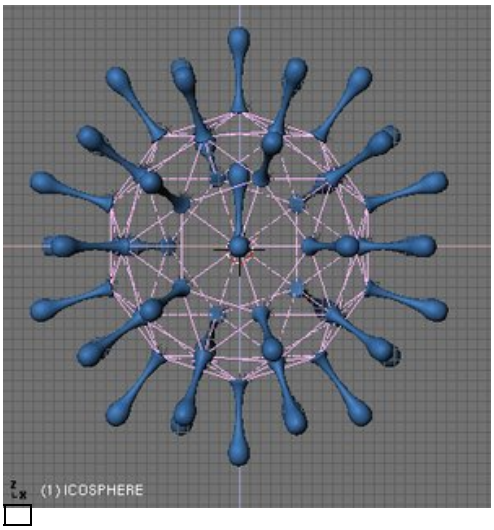
I had to take special care to be sure that the tentacle was located at the sphere center, and that both the tentacle axis and the sphere axis were aligned with the world axis as above (*Local reference of the tentacle*).

Now, I simply make the icosphere the parent of the tentacle. Select the icosphere alone and made it DupliVert in the Anim settings panel (*DupliVerts not rotated*).

Press the Rot button to rotate the tentacles (*DupliVerts rotated*).



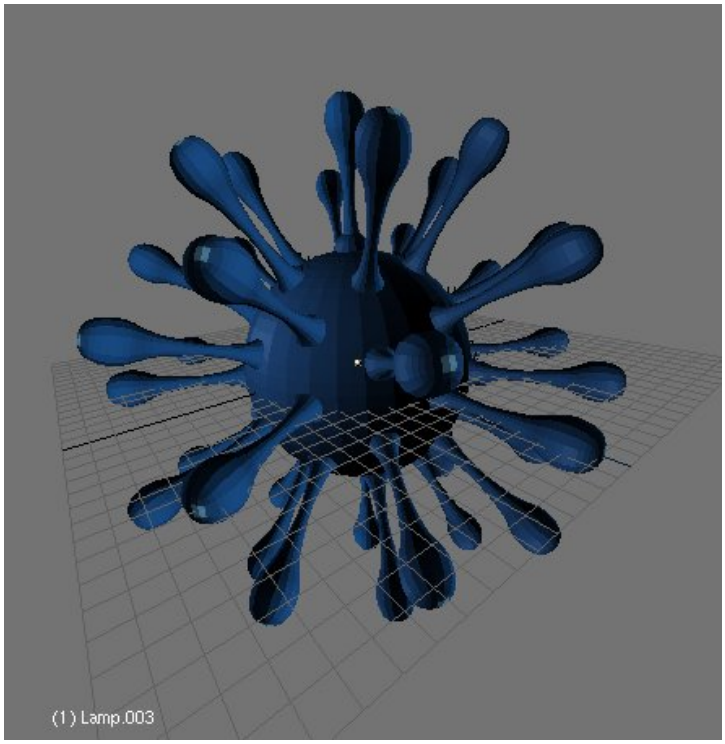
DupliVerts not rotated.



DupliVerts rotated.

Once again to make the tentacle point outwards we have to take a closer look to itâ← s axis. When applying Rot, Blender will try to align one of the tentacleâ← s axis with the normal vector at the parent mesh vertex.

We didnâ ← t care about the parent circle for Magritteâ ← s men, but here we should care about the sphere, and you will soon notice that it is *not* rendered. You probably would like to add an extra renderable sphere to complete the model. You can experiment in Edit mode with the tentacle, moving itâ ← s vertices off the center of the sphere, but the objectâ ← s center should always be at the sphereâ ← s center in order to get a symmetrical figure. However take care not to scale up or down in one axis in Object mode since it would lead to unpredictable results in the DupliVerted objects when applying the Rot button.



Our model complete.
Once youâ ← re done with the model and you are happy with the results, you can select the tentacle and press Ctrlâ § ShiftA and click on the Make duplis real ? menu to turn your virtual copies into real meshes (*Our model complete*).

See also

Other duplication methods are listed here.

DupliFaces

Mode: Object mode

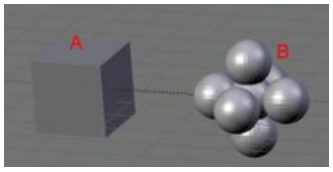
Panel: Anim settings â DupliFaces

Hotkey: F7

DupliFaces is the capability to replicate an object on each face of a parent object. One of the best ways to explain this is through an example illustration.

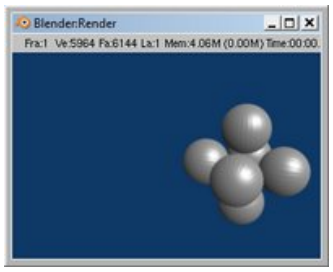


Base.
Lets start by examining the base for the story of DupliFaces. To use DupliFaces there must be two objects, one of which will be the parent of the other. In the diagram we have a cube called â← Aâ← that is a parent of the sphere labeled â← Bâ← .



Applied.
Next in the Object context (F7) â Anim settings panel, click on the DupliFaces button, which executes this effect. What happens is that the sphere which is a child of the cube is duplicated once for each face of the cube.

Generically, DupliFaces reproduces the child(ren) object(s) of a parent as many times as there are faces on this parent. The location, orientation and scale of the duplicated child(ren) matches that of the faces of the parent.



Rendered.
What is actually rendered is not exactly what is seen in the 3D view. The parent object is hidden from rendering and hence serves merely as a template. In addition, the original child(ren) object(s) is(are) also hidden so that only the duplicated instances of the object(s) are shown. In our example, the cube is not present in the render and neither is the original sphere which would otherwise have been in the centre of the six surrounding sphere corresponding to the six faces of the cube.

By default, the scaling of the duplicated objects maintains the scaling of the original child object. In the Anim settings panel there is a button called Scale that becomes visible when DupliFaces is selected. Pressing this button causes the duplicated objects to take on the scaling factor of the faces from which they were generated.

DupliGroup

Mode: Object mode

Panel: Anim settings â DupliGroup

Hotkey: F7

DupliGroup allows you to create an instance of a group with the transformation of the object.

Basic Usage

- Create a group of objects by selecting the objects to be grouped and pressing CtrlG â Add to new group.
- To name the group, select one of the objects in the group you just created and press F7 to switch to the Object context. Edit the name displayed in the GR: block of the Object and Links panel.
- Select Add â Group â `[name of group you just created]`.

DupliGroup and Dynamic Linking

See Appending and Linking to understand how to dynamically link data from another .blend file into the current file. You can dynamically link groups from one blend file to another. When you do so, the linked group does not appear anywhere in your scene until you create an object controlling where the group instance appears.

Example

- Link a group from another file into your scene, as described in Appending and Linking.

From here, you can use the easy way or the hard way:

- The easy way:

 1. Select Add â Group â `[name of group you just linked]`.

- The hard way:

 1. Select Add â Empty, and select the empty that you added.
 2. Switch to the Object context with F7.
 3. Click DupliGroup.
 4. Under GR:, type the name of the group that you linked.

At this point, an instance of the group will appear. You can duplicate the empty, and the DupliGroup settings will be preserved for each empty. This way, you can get multiple copies of linked data very easily.

Making a DupliGroup Objects Real

Say you want to make further edits on an DupliGroup instance or render the DupliGroup in Yafaray or some other render that does not support importing DupliGroups directly:

Simply select your DupliGroup and press Ctrlâ § ShiftA to convert the DupliGroup into an regular objects that can be transformed and animated normally.

Note that if the DupliGroup was linked from an external file the Object Data (mesh, materials, textures, transforms)will also still be linked from that file. However, the various object's parent-child relationships do not carry over.

Armatures

Often youâ ← ll want to animate the linked group â ← but by default, you canâ ← t move any object in the group individually! If you use DupliGroups and proxy objects, you can create multiple instances of a linked group. Development of this feature is a work in progress; in Blender 2.43 and CVS (as of 29 April 2007), a proxy object controls *all instances of a group*. It is not yet possible to have one proxy per group instance.

To animate with armatures:

- Link a group containing an armature from another file into your scene, as described in Appending and Linking.
- Select Add â Group â `[name of group you just linked]`.
- In the 3D window, select the empty that controls the new group instance and press CtrlAltP. Select the armature and youâ ← ll have a proxy for it.

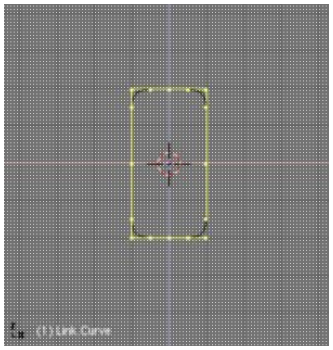
If you are using a POSIX compliant file system, you can work around the one proxy object per group limitation with the cheap hack documented at Linked Lib Animation Madness.

DupliFrames

You can consider DupliFrames in two different ways: an arranging or a modeling tool. In a way, DupliFrames are quite similar to DupliVerts. The only difference is that with DupliFrames we arrange our objects by making them follow a curve rather than using the vertex of a mesh. DupliFrames stands for â← *Dupli*cation at *Frames*â← and is a very useful modeling technique for objects which are repeated along a path, such as the wooden sleepers in a railroad, the boards in a fence or the links in a chain, but also for modeling complex curve objects like corkscrews, seashells and spirals.

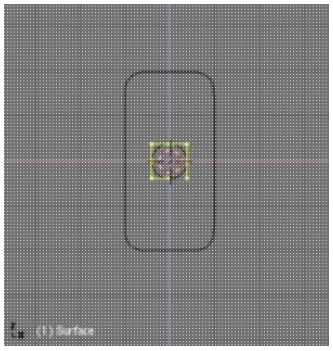
Modeling using DupliFrames

We are going to model a chain with its links using DupliFrames. First things come first. To explain the use of DupliFrames as a modeling technique, we will start by modeling a single link.



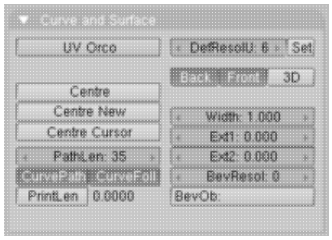
☐
Linkâ← s outline.
To do this, add in front view a Curve Circle (BÃ©zier or NURBS, whatever).

In Edit mode, subdivide it once and move the vertices a little to fit the linkâ← s outline (*Linkâ← s outline*).



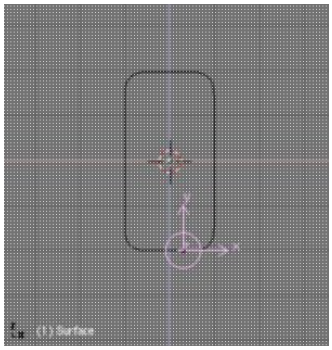
☐
Linkâ← s cross section.
Leave Edit mode and add a Surface Circle object (*Linkâ← s cross section*).

NURBS-surfaces are ideal for this purpose, because we can change the resolution easily after creation, and if we need to, we can convert them to a mesh object. It is very important that you do not confuse Curve Circle and Surface Circle. The first one will act as the shape of the link but it will not let us do the skinning step later on. The second one will act as a cross section of our skinning.

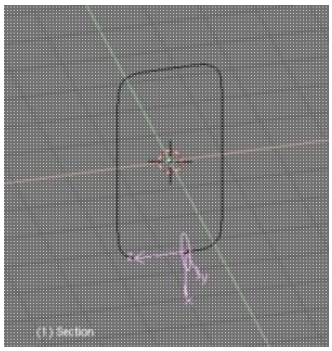


Curveâ← s settings: Curve Path and Curve Follow.
Now parent the circle surface to the circle curve (the linkâ← s outline) as a Normal Parent (not a Curve Follow constraint).

Select the curve and in the Editing context (F9), Curve and Surface panel, press Curve Path and Curve Follow (*Curveâ← s settings: Curve Path and Curve Follow*).

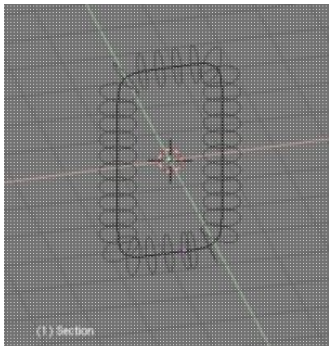


Clearing origin.
Itâ← s probable that the circle surface will appear dislocated. If this is the case, select it and press AltO to clear the origin (*Clearing origin*).

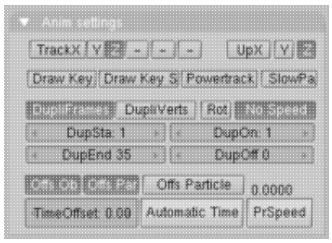


Tracking the right axis.
If you hit AltA the circle will follow the curve.

Now you probably will have to adjust the TrackX/Y/Z and UpX/Y/Z tracking buttons, in the Anim settings panel, Object context (F7), to make the circle go perpendicular to the curve path (*Tracking the right axis*).

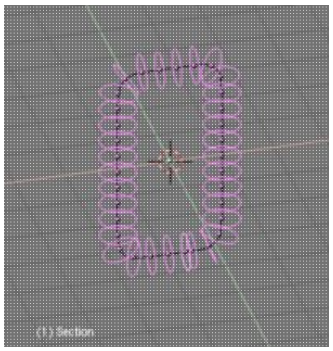


DupliFrames!
Now select the Surface Circle and, still in the Anim settings panel, press DupliFrames. A number of instances of the circular cross section will appear along the curve path (*DupliFrames!*).



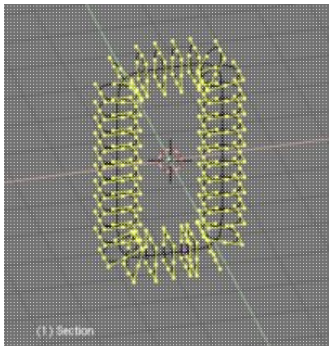
Values for DupliFrames. Note â← DupEnd: 35â← will end link before curveâ← s end.
You can adjust the number of circles you want to have with the DupSta, DupEnd, DupOn and DupOff buttons. These buttons control the Start and End of the duplication, the number of duplicates each time and also the Offset between duplications. If you want the link to be opened, you can try a different setting for DupEnd (*Values for DupliFrames. Note â← DupEnd: 35â← will end link before curveâ← s end*).

Note that the maximum number of duplications (the DupEnd value at which all curve path is â← usedâ←) is controlled by the â← lengthâ← of this curve path (Path Len field, in the Curve and Surface panel).



Making Dupliâ← s Real.
To turn the structure into a real NURBS-object, select the Surface Circle and press Ctrlâ § ShiftA.

A pop-up menu will appear, clic on Make dupli objects real (*Making Dupliâ← s Real*).



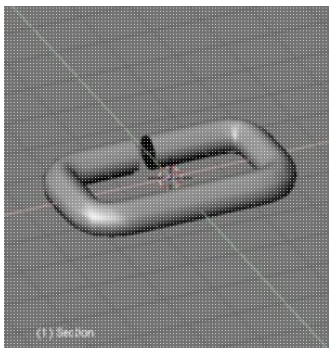
Skinning the link.
Do not deselect anything. We now have a collection of NURBS forming the outline of our object, but so far they are not skinned, so we cannot see them in a shaded preview or in a rendering.

To achieve this, we need to join all the rings to one object. Without deselecting any rings, press CtrlJ and confirm the pop-up menu request.

Now, enter Edit mode for the newly created object and press A to select all vertices (*Skinning the link*).

Now we are ready to skin our object.

Press F and Blender will automatically generate the solid object. This operation is called *Skinning* and is fully described in â← Surfaces Skinningâ← .



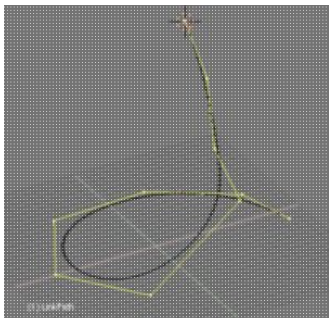
Skinned link.
When you leave Edit mode, you can now see the object in a shaded view.In some versions of Blender it may appear very dark. To correct this, enter Edit mode to flip all normals: select all vertices, then press W. Choose Flip Normals from the menu and leave Edit mode. The object will now be drawn correctly (*Skinned link*).

The object we have created is a NURBS object. This means that you can still edit it. Even more interestingly, you can also control the resolution of the NURBS object via the settings in the Curve and Surface panel, Editing context (F9).

Here you can set the resolution of the object using ResolU and ResolV, so you can adjust it for working with the object in a low resolution, and then set it to a high resolution for your final render. NURBS objects are also very small in file size for saved scenes. Compare the size of a NURBS scene with the same scene in which all NURBS are converted (AltC) to meshes.

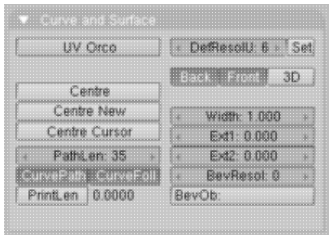
Finally you can delete the curve we used to give the shape of the link, since we no longer need it.

Arranging objects with DupliFrames

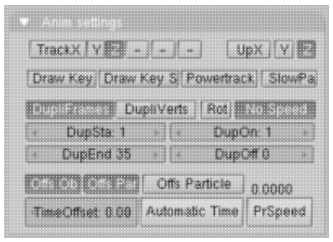


Using a curve path to model the chain.
Now we will continue modeling the chain itself. For this, just add a Curve Path (we could use a different curve but this one gives better results). In Edit mode, move its vertices until you get the desired shape of the chain (*Using a curve path to model the chain*).

If not using a Curve Path, you should check the button 3D (in the Editing context, Curve and Surface panel), to let the chain be real 3D.

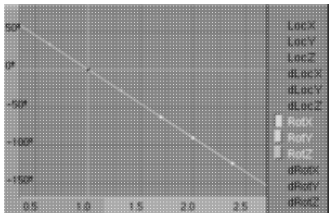


Curve settings.
Select the object â← Linkâ← we modeled in the previous step and parent it to the chain curve, again as a Normal Parent. Since we are using a Curve Path the option CurvePath, in the Editing context, will be automatically activated, however the CurveFollow option will not, so you will have to activate it (*Curve settings*).

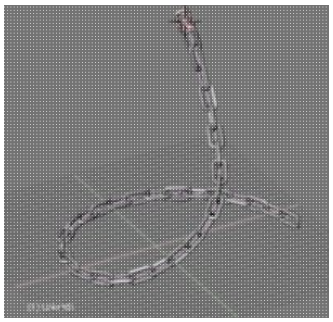


Adjusting the DupliFrames.
If the link is dislocated, select it and press AltO to clear the origin. Until now we have done little more than animate the link along the curve. This can be verified by playing the animation with AltA.

Now, with the link selected once again go to the Object context (F7), Anim settings panel. Here, activate the option DupliFrames as before. Play with the DupSta, DupEnd and DupOff NumButtons. Normally we are going to set DupOff to 0, but for a chain, if this produce links too close from each other, you should change the value PathLen for the path curve to a lesser value (Editing context, Curve and Surface panel) and then correspondingly change the DupEnd value for the link to that same number (*Adjusting the DupliFrames*).



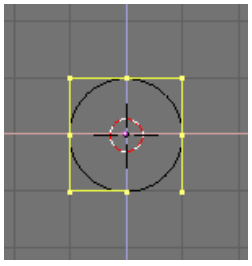
Rotating the link.
We need it so that the link rotates along the curve animation, so we have each link rotated 90 degrees with respect to the preceding one in the chain. For this, select the link and press Axis in the Draw panel, Object context, to reveal the objectâ ← s axis. Insert a rotation keyframe in the axis which was parallel to the curve. Move 3 or 4 frames ahead and rotate along that axis pressing R followed by X-X (X twice), Y-Y, or Z-Z to rotate it in the *local* X, Y or Z axis (*Rotating the link*).



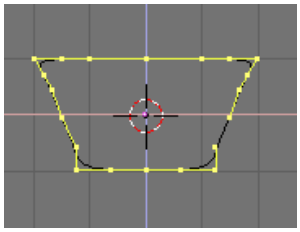
Dupliframed chain.
Open an Ipo window to edit the rotation of the link along the path. Press the Extrapolation Mode (E â Extrapolation) so the link will continually rotate until the end of the path. You can edit the Ipo rotation curve to make the link rotate exactly 90 degrees every one, two or three links (each link is a frame). Use N to locate a node exactly at X=2.0 and Y=9.0, which correspond to 90 degrees in 1 frame (from frame 1 to 2). Now we got a nice chain (*Dupliframed chain*)!

More Animation and Modeling

You are not limited to use Curve Paths to model your stuff. These were used just for our own convenience, however in some cases there is no need of them.

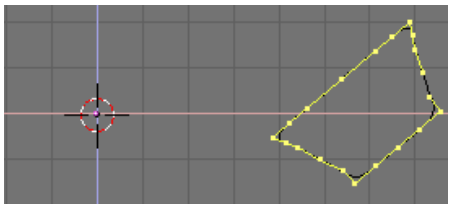


A Surface Circle.
In Front View, add a surface circle (you should know why by now, *A Surface Circle*).



Trapezoidal cross-section.
Subdivide once, to make it look more like a square.

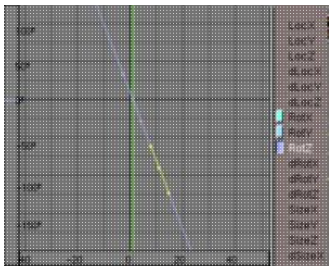
Move and scale some vertices a little to give it a trapezoid shape (*Trapezoidal cross-section*).



Trapezoidal cross section, rotated and translated.
Then rotate all vertices a few degrees.

Grab all vertices and displace them some units right or left in X (but at the same Z location). You can use Ctrl to achieve this precisely.

Leave Edit mode (*Trapezoidal cross section, rotated and translated*).



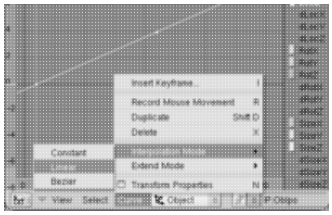
□
Rotation Ipo for the cross section.
From now on, the only thing we are going to do is editing Ipo animation curves. So you can call this â← Modeling with Animationâ← if you like. We will not enter Edit mode for the surface any more.

Switch to Top View.

Insert a KeyFrame for rotation at frame 1, go ahead 10 frames and rotate the surface 90 degrees over its new origin.

Insert one more KeyFrame.

Open an Ipo window, and set the rotation Ipo to Extrapolation mode (E â Extrapolation â← *Rotation Ipo for the cross section*).



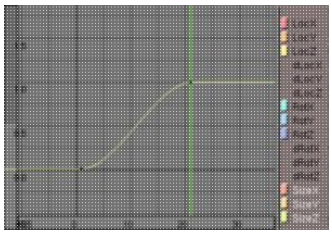
Translation Ipo for the cross section.
Go back to frame 1 and insert a keyframe for Location.

Switch to Front View.

Go to frame 11 (just press â) and move the surface in Z a few grid units.

Insert a new keyframe for Location.

In the Ipo window set the LocZ to Extrapolation mode (*Translation Ipo for the cross section*).



Size Ipo for the cross section.
Now, of course, go to the Object context and press DupliFrames. You can see how our surface is ascending in a spiral through the 3D space forming something like a spring. This is nice, however we want more.
Deactivate DupliFrames to continue.

In frame 1 scale the surface to nearly zero and insert a keyframe for Size.

Go ahead to frame 41, and clear the size with AltS.

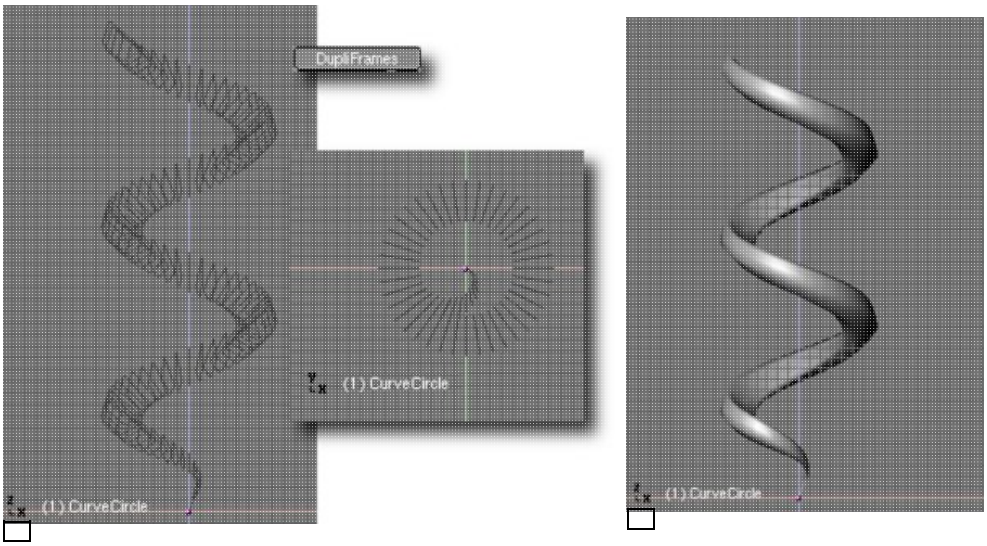
Insert a new keyframe for size.

This Ipo will not be in extrapolation mode since we donâ← t want it scaled up at infinitum (*Size Ipo for the cross section*).

If you now activate DupliFrames you will see a beautiful outline of a corkscrew (*A curved object procedurally created*).

Once again the last steps are:

- Make Duplis Real.
- Joining the surfaces.
- Select all vertices and skinning.
- Switch direction of normal if needed.
- Leave Edit mode (*A curved object procedurally created*).



A curved object procedurally created.
You can see this was a rather simple example. With more Ipo curve editing you can achieve very interesting and complex models. Just use your imagination.

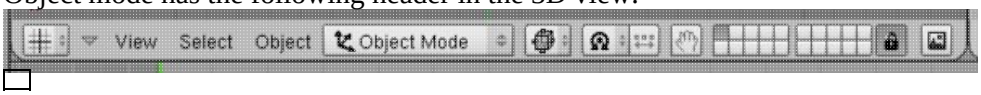
Edit Mode

How to enter Edit mode

You can work with geometric objects in two modes.

Object mode

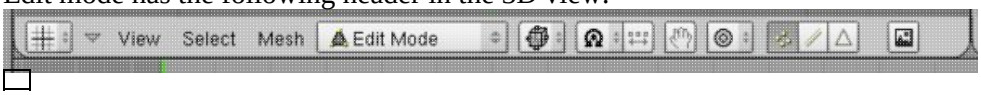
Operations in Object mode affect whole object.
Object mode has the following header in the 3D view:



Object mode header.

Edit mode

Operations in Edit mode affect only the geometry of an object, but not global properties such as location or rotation.
Edit mode has the following header in the 3D view:



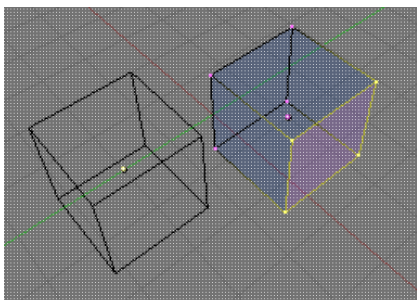
Edit mode header.

You switch between these two modes with the â Tab key or by selecting the desired Mode in the menu in the 3d view header.

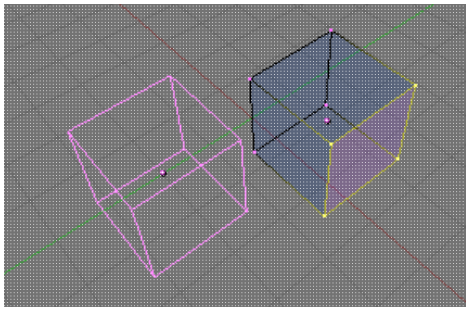
After creating an object you may be immediately placed in Edit mode â← depending on whether the Switch to Edit Mode button is toggled, in the User Preferences window, Edit Methods â← tabâ← , Add new objects: group. Edit mode only applies to one object at a time, the *active*, or most recently selected, object.

Visualization

An object shown in wireframe when selected in Object mode) is drawn in purple in 3D View windows; it is (generally!) black when unselected.



One cube selected.
In Edit mode, each unselected vertex is drawn in purple, each edge is drawn in black and each face is drawn in translucent dark-blue. Each selected vertex or edge is highlighted in yellow. In the image (*One cube selected*) the cube on the right is in Edit mode. The cube on the left is in Object mode *and* not selected.



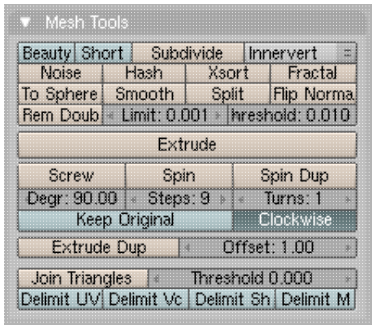
Two cubes selected prior to Edit mode.

When Edit mode is chosen with multiple objects selected only the most recently selected or *active* object) enters Edit mode. The other objects remain purple and in Object mode. As shown in the image (*Two cubes selected prior to Edit mode*), both cubes were selected prior to Edit mode and now the left cube is still purple and only the right cube (the active one) is in Edit mode.

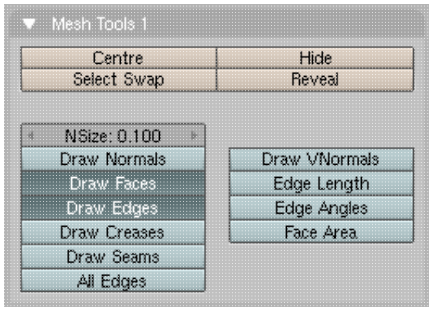
If enough vertices are selected to form a face then that face is highlighted in translucent purple while the remaining faces are highlighted in translucent dark-blue. This helps give you a frame of reference when selecting vertices, edges or faces. The translucent effect indicates that you have selected enough vertices to imply one or more faces. See â← Edge and Face Toolsâ← for further details on implicit selections.

Mesh panels

If the Buttons window is visible and Editing context button (F9) is activated, then two panels appear while in Edit mode (see images *Mesh Tools* and *Mesh Tools 1*):



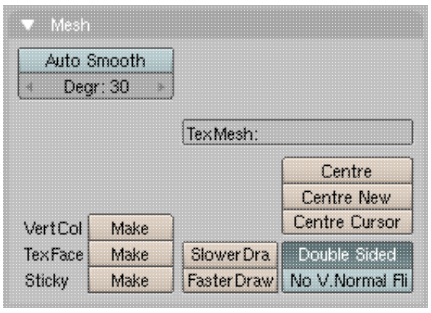
Mesh Tools.



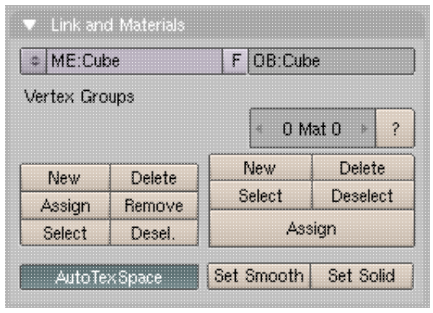
Mesh Tools 1.

By default the buttons Draw Faces and Draw Edges are pre-selected and any selected edges and faces are highlighted.

In addition, panels (see images *Link and Materials* and *Mesh*) are updated.



Mesh.

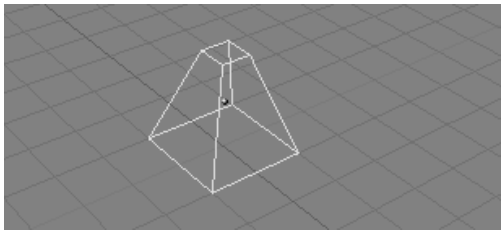


Link and Materials.

The Link and Materials panel gains the Vertex Groups buttons group (New, Delete, Assign, Remove, Select and Desel.). The Mesh panel loses the Decimator, Apply and Cancel group of buttons.

Basic Editing

Most simple operations from Object mode (like selecting, moving, rotating, and scaling) work the same way on vertices as they do on objects. Thus, you can learn how to handle basic Edit mode operations very quickly. The only notable difference is a new scaling option, AltS, which scales the selected vertices along the direction of the normals (shrinks-fattens). The truncated pyramid in image (*Chopped-off pyramid*), for example, was created with the following steps:



Chopped-off pyramid.

1. Add a cube to an empty scene. If not in Edit mode then use â Tab to enter Edit mode.
2. Make sure all vertices are deselected (purple). Use Border select to select the upper four vertices.
3. Check that the scaling center is set to *anything but* the 3D Cursor (you *do not* want to see as the selected pivot point), then switch to scale mode (S), reduce the size, and confirm with LMB .
4. Exit Edit mode by pressing â Tab.

All operations in Edit mode are ultimately performed on the vertices; the connected edges and faces automatically adapt, as they depend on the verticesâ← positions. To select an edge, you must select the two endpoints or place the mouse on the edge and press AltRMB . To select a face, each corner must be selected.

Edit mode operations are many, and most are summarised in the Editing context of the Buttons window, accessed via the () header button or via F9.

Mirror Axis and Modifier



Mirror Axis.
One extra feature for Edit mode is the mirroring tool. If you have some vertices selected and you press M you will be presented with a menu containing nine options (*Mirror Axis*). You can select from these to mirror the selected vertices with respect to any of the X, Y or Z axes of the Global, Local, or View space reference. If you need to select groups of vertices use the handy Loop to Region tool.

Editorâ← s Note

There is a much more advanced tool for performing mirroring operations and that is the Mirror modifierâ←¦

Specials

With W you can call up the Specials menu in Edit mode, see image (*Specials menu*). With this menu you can quickly access functions which are frequently required for polygon-modeling.



Specials menu.

- Subdivide â← Each selected edge is split in two, new vertices are created at middle points, and faces are split too, if necessary.
- Subdivide Multi â← This is identical to Subdivide except a dialog pops up asking for the *number of cuts* or repeated sub-divisioning. The default is â← **2**â← .
- Subdivide Multi Fractal â← As above, but new vertices are randomly displaced within a user-defined range.
- Subdivide Smooth â← Same as Subdivide, but new vertices are displaced towards the barycenter (center of mass) of the connected vertices.
- Merge â← Merges selected vertices into a single one, at the barycenter position or at the 3D cursor position.
- Remove Doubles â← Merges all of the selected vertices whose relative distance is below a given threshold (0.001 by default).
- Hide â← Hides selected vertices (H).
- Reveal â← Shows hidden vertices (AltH).
- Select Swap â← All selected vertices become unselected and vice-versa (CtrlI).
- Flip Normals â← Change the normal directions of the selected faces.
- Smooth â← Smooths out a mesh by moving each vertex towards the barycenter of the linked vertices.
- Bevel â← Bevels the entire object regardless of the selected vertices, edges or faces. See this page.
- Set Smooth â← Changes the selected faces to smoothing shading.
- Set Solid â← Changes the selected faces to faceted or flat shading.



Keyboard Tip:

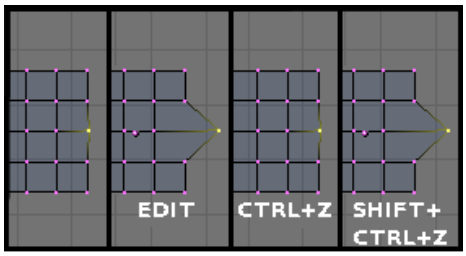
You can access the entries in a popup-menu by using the corresponding numberkey. For example, pressing W and then 1 will subdivide the selected edges without you having to touch the mouse at all.



Mesh Tools panel.
Many of these actions have a button of their own in the Mesh Tools panel of the Buttons window, Editing context, see image (*Mesh Tools panel*). The Remove doubles threshold can be adjusted on that panel too.

Mesh Undo/Redo

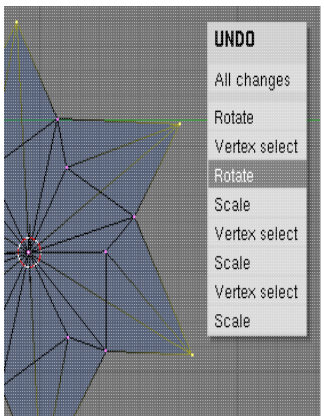
Blender has a global undo system, giving full multiple level undo capabilities in all areas of Blender. Exceptions are: File Browser, Audio and Outliner windows.



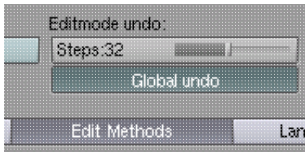
Undo and Redo.
The new global hotkey for undo is CtrlZ, and CtrlY for redo.

Mesh undo works in the background saving copies of your mesh in memory as you make changes. Pressing the CtrlZ in mesh Edit mode reverts to the previously saved mesh, undoing the last edit operation (*Undo and Redo*).

Mesh â← Undoâ← operations are only stored for one mesh at a time. You can leave and re-enter Edit mode for the same mesh without losing any undo information, but once another mesh is edited, the undo information for the first is lost. Pressing CtrlY (or Ctrlâ § ShiftZ) re-does the last undo operation (*Undo and Redo*).



Editmode Undo History menu.
Pressing AltU brings up the Editmode Undo History menu, see image (*Editmode Undo History menu*). This lists all the undo steps by name so you can quickly find your way back to a known good point in your work. The â← undoâ← menu do not have an option to reload the mesh data as it was at the beginning of your edit session, but you can achieve the same effect by returning in Object mode (â Tab), hit AltU to show the Global Undo History menu, and select the step just before the edit session you want to undo.

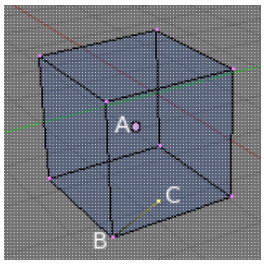


User Preferences â Edit Methods.
Mesh â← Undoâ← has the potential to be very memory intensive. A mesh of 64,000 faces and vertices can use over 3MBs of RAM per undo step! If you are on a machine that is strapped for RAM (memory), in the User Preferences window under Edit Methods, there is a NumButton for setting the maximum number of undo steps saved, see image (*User Preferences â Edit Methods*). The allowable range is between 1 and 64. The default is 32.

Vertices, Edges and Faces

In basic meshes, everything is built from three basic structures: *Vertices*, *Edges* and *Faces* (weâ← re not talking about curves, NURBS, and so forth here). But there is no need to be disappointed: this simplicity still provides us with a wealth of possibilities that will be the foundation for all our models.

Vertices



Vertex example.
A vertex is primarily a single point or position in 3D space. It is usually invisible in rendering and in Object mode. Donâ← t mistake the centre point of an object for a vertex. It looks similar, but itâ← s bigger and you canâ← t select it. (*Vertex example*) shows the centre point labelled as â← Aâ← . â← Bâ← and â← Câ← are vertices.

To create a new vertex, change to Edit mode, and click CtrlLMB . Of course, as a computer screen is two-dimensional, Blender canâ← t determine all three vertex coordinates from one mouse click, so the new vertex is placed at the depth of the 3D cursor â← intoâ← the screen. Any vertices selected previously are automatically connected to the new one with an edge. Vertex labelled â← Câ← is a new vertex added to the cube with a new edge (B to C).

Edges

An edge always connects two vertices with a straight line. The edges are the â← wiresâ← you see when you look at a mesh in wireframe view. They are usually invisible on the rendered image. They are used to construct faces. Create an edge by selecting two vertices and pressing F.

Faces

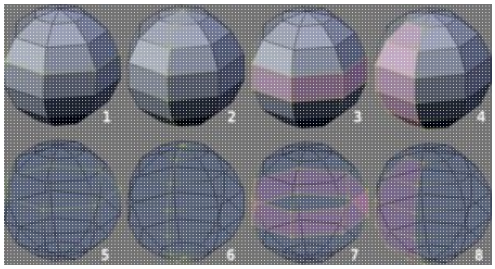
A face is the highest level structure in a mesh. Faces are used to build the actual surface of the object. They are what you see when you render the mesh. A face is defined as the area between either three (triangles) or four (quadrangles) vertices, with an edge on every side. Triangles always work well, because they are always flat and easy to calculate.

Take care when using four-sided faces (quads), because internally they are simply divided into two triangles each. Four-sided faces only work well if the face is pretty much flat (all points lie within one imaginary plane) and convex (the angle at no corner is greater than or equal to 180 degrees). This is the case with the faces of a cube, for example. Thatâ← s why you canâ← t see any diagonal in its wireframe model, because they would divide each square face into two triangles.

While you could build a cube with triangular faces, it would just look more confusing in Edit mode. An area between three or four vertices, outlined by edges, doesnâ← t have to be a face. If this area does not contain a

face, it will simply be transparent or non-existent in the rendered image. To create a face, select three or four suitable vertices and press F.

Edge Loops and Face Loops



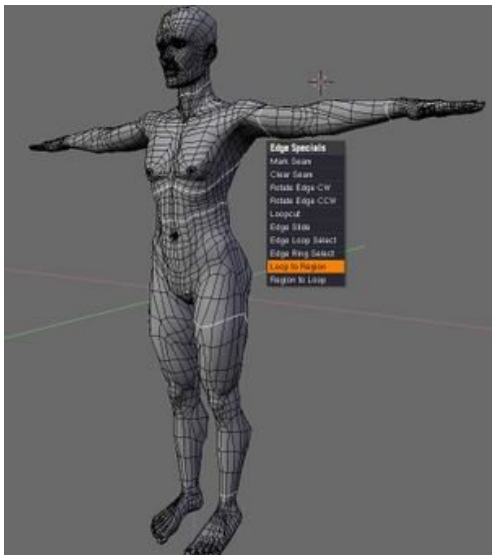
Edge and Face Loops.
Edge and Face Loops are sets of faces or edges that form continuous â← loopsâ← as shown in (*Edge and Face Loops*). The top row (1-4) shows a solid view, the bottom row (5-8) a wireframe view of the same loops.

Note that loops 2 and 4 do not go around the whole model. Loops stop at so called poles because there is no unique way to continue a loop from a pole. Poles are vertices that are connected to either three or five or more edges. Accordingly, vertices connected to exactly one, two or four edges are not poles.

Loops that do not end in poles are cyclic (1 and 3). They start and end at the same vertex and divide the model into two partitions. Loops can be a quick and powerful tool to work with specific, continuous regions of a mesh and are a prerequisite for organic character animation. For a detailed description of how to work with loops in Blender, please refer to the Manual page on Edge and Face Tools.

Edge Loops

Loops 1 and 2 in (*Edge and Face Loops*) are edge Loops. They connect vertices so that each one on the loop has exactly two neighbours that are not on the loop and placed on both sides of the loop (except the start and end vertex in case of poles).



Edge Loops in organic modeling.
Edge Loops are an important concept especially in organic (subsurface) modeling and character animation. When used correctly, they allow you to build models with relatively few vertices that look very natural when used as subdivision surfaces and deform very well in animation.

Take (*Edge Loops in organic modeling*) as an example: the edge loops follow the natural contours and deformation lines of the skin and the underlying muscles and are more dense in areas that deform more when the character moves, for example at the shoulders or knees.

Further details on working with Edge Loops can be found in Edge Loop Selection.

Face Loops

These are a logical extension of Edge Loops in that they consist of the faces between two Edge Loops, as shown in loops 3 and 4 in (*Edge and Face Loops*). Note that for non-circular loops (4) the faces containing the poles are not included in a Face Loop.

Further details on working with Face Loops can be found in Face Loop Selection.

Basic Mesh Objects

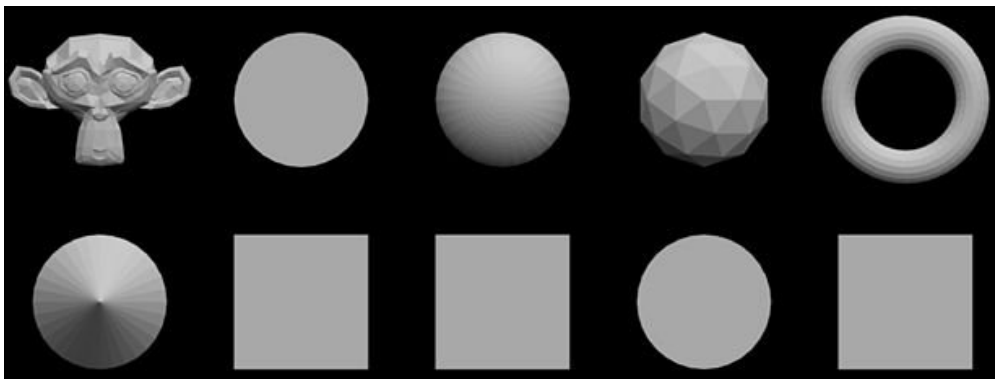
Mode: Object mode

Hotkey: â § ShiftA

Menu: Add Â» Mesh

Description

A common object type used in a 3D scene is a mesh. Blender comes with a number of â← primitiveâ← mesh shapes that you can start modeling from.



All the mesh primitives that can be created (none smoothed).
Options included in more than one primitive are:

Radius
: Sets the starting size for Circle, Cylinder, Cone, UVSphere and IcoSphere.

Depth
: Sets the starting length for Cylinder and Cone.

Note about planar primitives
You can make a planar mesh three-dimensional by moving one or more of the vertices out of its plane (applies to Plane, Circle and Grid).

Options

Plane
: A standard plane contains four vertices, four edges, and one face. It is like a piece of paper lying on a table; it is not a real three-dimensional object because it is flat and has no thickness. Objects that can be created with planes include floors, tabletops, or mirrors.

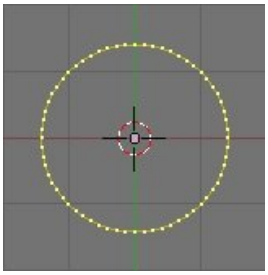
Cube
: A standard cube contains eight vertices, twelve edges, and six faces, and is a real three-dimensional object. Objects that can be created out of cubes include dice, boxes, or crates.



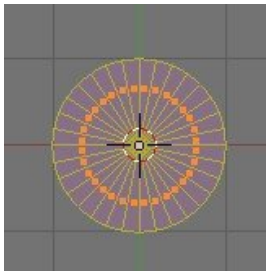
Add Circle popup window.

Circle

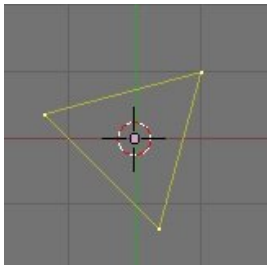
A standard circle is comprised of *n* vertices. The number of vertices and radius can be specified in the popup window which appears when the circle is created, shown in (*Add Circle popup window*). When Fill button is active, the circle will be filled with triangular faces which share a vertex in the middle. However the circle is only a flat shape. If it is not filled and you want to render it, you must assign it a wireframe material (Shading context â← F5 â← , Material sub-context, Links and Pipeline panel and finally the Wire button). The Radius parameter adjusts the size of the circle.
The more vertices the circle contains, the smoother its contour will be, see (*â← Circlesâ← obtained with various settings*).



A circle with 64 vertices gives a smooth circle.



Filled circle with 32 vertices.



A circle with only 3 vertices is actually a triangle.



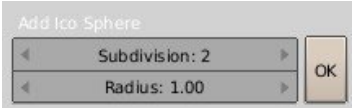
Add UV Sphere popup window.

UVSphere

A standard UVsphere is made out of *n* segments and *m* rings. The level of detail and radius can be specified in the popup window which appears when the UVsphere is created. Increasing the number of segments and rings makes the surface of the UVsphere smoother. Segments are like Earthâ← s meridians, going pole to pole and Rings are like Earthâ← s parallels. Example objects that can be created out of UVspheres are balls, heads or pearls for a necklace.

Note

If you specify a six segment, six ring UVsphere youâ← ll get something which, in top view, is a hexagon (six segments), with five rings plus two points at the poles. Thus, one ring fewer than expected, or one more, if you count the poles as rings of radius 0.



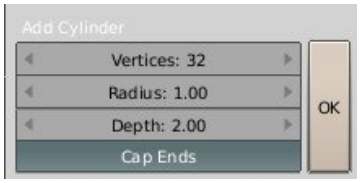
Add Ico Sphere popup window.

Icosphere

An Icosphere is made up of triangles. The number of subdivisions and radius can be specified in the

window that pops up when the Icosphere is created; increasing the number of subdivisions makes the surface of the Icosphere smoother. At level 1 the Icosphere is an icosahedron, a solid with 20 equilateral triangular faces. Any increasing level of subdivision splits each triangular face into four triangles, resulting in a more spherical appearance. Icospheres are normally used to achieve a more isotropical and economical layout of vertices than a UVsphere.

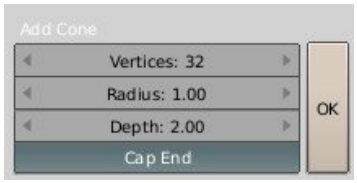
Note
It is possible to add an icosphere subdivided 500 times. Adding such a dense mesh is a sure way to cause a program crash. An icosphere subdivided 10 times would have 5,242,880 triangles, so be very careful about this!



Add Cylinder popup window.

Cylinder

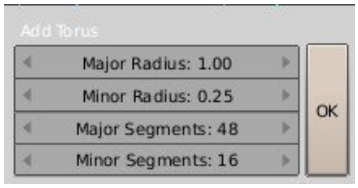
A standard cylinder is made out of *n* vertices. The number of vertices in the circular cross-section can be specified in the popup window that appears when the object is created; the higher the number of vertices, the smoother the circular cross-section becomes. The Radius and Depth parameters controls dimensions of cylinder. Objects that can be created out of cylinders include handles or rods.
If Cap Ends is inactive, the created object will be a tube. Objects that can be created out of tubes include pipes or drinking glasses (the basic difference between a cylinder and a tube is that the former has closed ends).



Add Cone popup window.

Cone

A standard cone is made out of *n* vertices. The number of vertices in the circular base, dimensions and option to close the base of cone can be specified in the popup window that appears when the object is created; the higher the number of vertices, the smoother the circular base becomes. Objects that can be created out of cones include spikes or pointed hats.



Add Torus popup window.

Torus

A doughnut-shaped primitive created by rotating a circle around an axis. The overall dimensions are defined by the Major and Minor Radius. The number of vertices (in segments) can be different for the circles and is specified in the popup window with both radii (Major Segments and Minor Segments).

Grid

A standard grid is made out of *n* by *m* vertices. The resolution of the x-axis and y-axis can be specified in the popup window which appears when the object is created; the higher the resolution, the more vertices are created. Example objects that can be created out of grids include landscapes (with the proportional editing tool or Displace modifier) and other organic surfaces. You can also obtain a grid when you create a plane and then use a subdivide modifier in Edit mode.

Monkey

This is a gift from old NaN to the community and is seen as a programmerâ← s joke or â← Easter Eggâ← . It creates a monkeyâ← s head once you press the Monkey button. The Monkeyâ← s name is â← Suzanneâ← and is Blenderâ← s mascot. Suzanne is very useful as a standard test mesh, much like the Utah Tea Pot or the Stanford Bunny.

Mesh smoothing

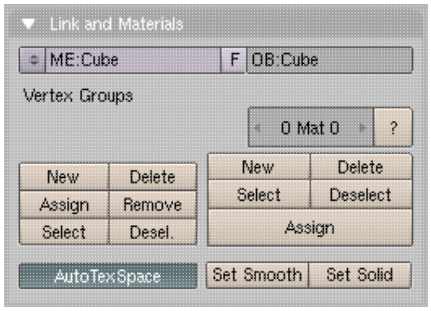


Simple un-smoothed test object.
As seen in the previous sections, polygons are central to Blender. Most objects are represented by polygons and truly curved objects are often approximated by polygon meshes. When rendering images, you may notice that these polygons appear as a series of small, flat faces. See (*Simple un-smoothed test object*).

Sometimes this is a desirable effect, but usually we want our objects to look nice and smooth. This section shows you how to visually smooth an object, and how to apply the Auto Smooth filter to quickly and easily combine smooth and faceted polygons in the same object.

The last section shows the possibilities to smooth a meshâ← s geometry, not only its appearance.

Smoothing the entire mesh



Link and Materials panel.
The easiest way is to set an entire object as smooth or faceted by selecting a mesh object, in Object mode, switching to the Editing context (F9), and clicking the Set Smooth button in the Link and Materials panel shown in (*Link and Materials panel*).

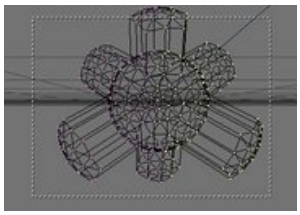
The button does not stay pressed, it forces the assignment of the â← smoothingâ← attribute to each face in the mesh, also when you add or delete geometry.



Completely smoothed.
Now, rendering the image with F12 should produce the image shown in (*Completely smoothed*).

Notice that the outline of the object is still strongly faceted. Activating the smoothing features doesnâ← t actually modify the objectâ← s geometry; it changes the way the shading is calculated across the surfaces, giving the illusion of a smooth surface. Click the Set Solid button in the same panel to revert the shading back to that shown in (*Simple un-smoothed test object*) above.

Manually smoothing parts of a mesh

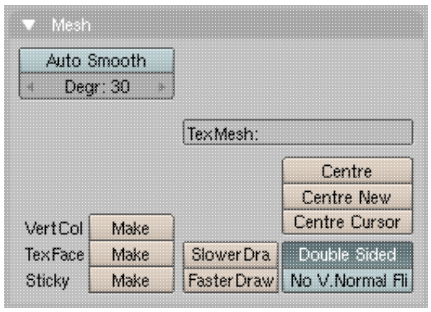


Object in Edit mode with some faces selected.
Alternatively, you can choose which faces to smooth by entering Edit mode for the object with â Tab, then selecting the faces and clicking the Set Smooth button (*Object in Edit mode with some faces selected*). The selected faces are marked in yellow.

When the mesh is in Edit mode, only the selected faces will receive the â← smoothingâ← attribute. You can set solid faces (removing the â← smoothingâ← attribute) in the same way by selecting faces and clicking the Set Solid button.

Auto Smooth

It can be difficult to create certain combinations of smooth and solid faces using the above techniques alone. Though there are workarounds (such as splitting off sets of faces by selecting them and pressing Y), there is an easier way to combine smooth and solid faces, by using Auto Smooth.



Auto Smooth button group in the Editing context of the Buttons window.
Press the Auto Smooth button in the Mesh panel of the Editing context, see (*Auto Smooth button group in the Editing context, Buttons window*) and Set Smooth to indicate which faces should be smoothed on the basis of the angle between faces. Angles on the model that are smaller than the angle specified in the Degr NumButton will be smoothed during rendering. Higher values will produce smoother faces, while the lowest setting will look identical to a mesh that has been set completely solid.

Note that a mesh, or any faces that have been set as Solid, will not change their shading when Auto Smooth is activated: this allows you extra control over which faces will be smoothed and which ones wonâ← t by overriding the decisions made by the Auto Smooth algorithm.



Same test object with Auto Smooth enabled.
Look at (*Same test object with Auto Smooth enabled*): only faces that have been set as *smooth* have been affected by the Auto Smooth feature.

Smoothing the mesh geometry

The above techniques do not alter the mesh itself, only the way it is displayed and rendered. Instead of just making the mesh look like a smooth surface, you can also physically smooth the geometry of the mesh with these tools:

You can apply one of the following in Edit mode:

- Smooth
- Subdivide Smooth
- Bevel

Alternatively, you can smooth the mesh non-destructively with one of the following modifiers:

- Smooth

 Works like the Smooth tool in Edit mode; can be applied to specific parts of the mesh using vertex groups.

- Subdivision Surface

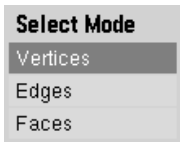
 Catmull-Clark subdivision produces smooth results. Sharp edges can be defined with subdivision creases or by setting certain edges to â← sharpâ← and adding an EdgeSplit modifier (set to From Marked As Sharp) before the Subsurf modifier.

Choosing which elements are selectable

There are many ways to select elements, and it depends on what mode you are in as to what selection tools are available. First we will go through these modes and after that a look is taken at basic selection tools.

Select Mode popup

In Edit mode there are three different selection modes, see (*Edit mode selection menu*).



Edit mode selection menu.

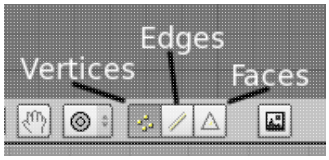
Select Mode Â» Vertices
: Press Ctrlâ Tab and select Vertices from the popup menu, or press Ctrlâ Tab1. The selected vertices are drawn in yellow and unselected vertices are drawn in a pink colour.

Select Mode Â» Edges
: Press Ctrlâ Tab and select Edges from the popup menu, or press Ctrlâ Tab2. In this mode the vertices are not drawn. Instead the selected edges are drawn in yellow and unselected edges are drawn in a black colour.

Select Mode Â» Faces
: Press Ctrlâ Tab and select Faces from the popup menu, or press Ctrlâ Tab3. In this mode the faces are drawn with a selection point in the middle which is used for selecting a face. Selected faces are drawn in yellow with the selection point in orange, unselected faces are drawn in black.

Almost all modification tools are available in all three modes. So you can Rotate, Scale, Extrude, etc. in all modes. Of course rotating and scaling a *single* vertex will not do anything useful, so some tools are more or less applicable in some modes.

Select Mode header widgets



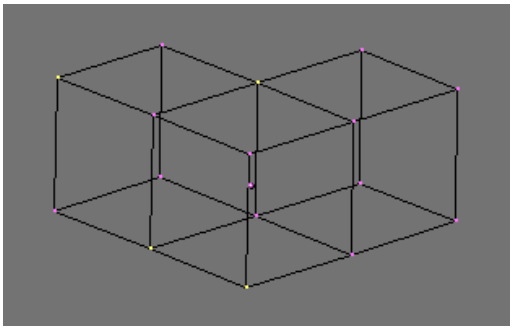
Edit mode selection buttons.
You can also enter the different modes by selecting one of the three buttons in the toolbar; see (*Edit mode selection buttons*).

Using the buttons you can also enter **mixed mode** by â § ShiftLMB clicking the buttons.

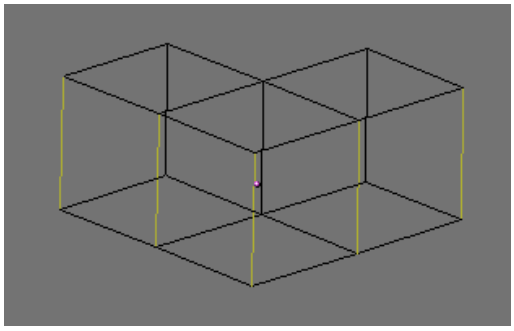
Note
The â← Mode Selectionâ← buttons are only visible in Edit mode.

Selected elements after switching selection mode

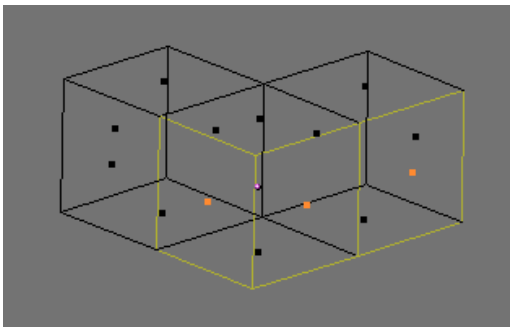
When switching modes, from Vertices to Edges and from Edges to Faces, the selected parts will still be selected if they form a complete set in the new mode. For example, if all four edges in a face are selected, switching from Edges mode to Faces mode will keep the face selected. All selected parts that do not form a complete set in the new mode will be unselected. See (*Vertices mode example*), (*Edges mode example*), (*Faces mode example*) and (*Mixed mode example*) for examples of the different modes.



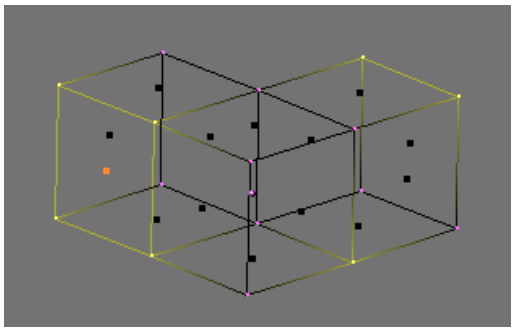
Vertices mode example.



Edges mode example.



Faces mode example.



Mixed mode example.

Basic selection

The most common way to select an element is to RMB  on that item, this will replace the existing selection with the new item.

Adding to a selection

To add to the existing selection hold down â § Shift while right clicking. Clicking again on a selected item will de-select it.

Selecting different types of faces

In Faces select mode, faces can be selected based on whether they are *triangles*, *quads*, or *others* (?). Hotkeys:

- CtrlAltâ § Shift3 selects all triangles.
- CtrlAltâ § Shift4 selects all quads.
- CtrlAltâ § Shift5 selects all non triangle/quad faces.

These tools are also available in both the 3D view header and toolbox Select menus.

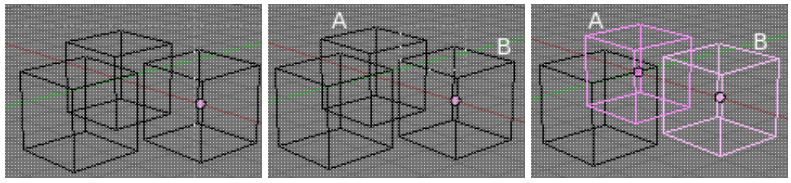
Selecting elements in a Region

Region selection allows you to select groups of elements within a 2D region. The region can be either a circle or rectangle. The circular region is only available in Edit mode. The rectangular region, or â← *Border Select*â← , is available in both Edit mode and Object mode.

Rectangular region (Border select)

Border Select is available in either Edit mode or Object mode. To activate the tool use the B. Use Border Select to select a group of objects by drawing a rectangle while holding down LMB. In doing this you will select all objects that lie within or touch this rectangle. If any object that was last active appears in the group it will become selected *and* active.

In (*Start*), Border Select has been activated and is indicated by showing a dotted cross-hair cursor. In (*Selecting*), the *selection region* is being chosen by drawing a rectangle with the LMB. The rectangle is only covering cubes â← Aâ← and â← Bâ← . Finally, by releasing LMB the selection is complete; see (*Complete*).



Start. Selecting. Complete.

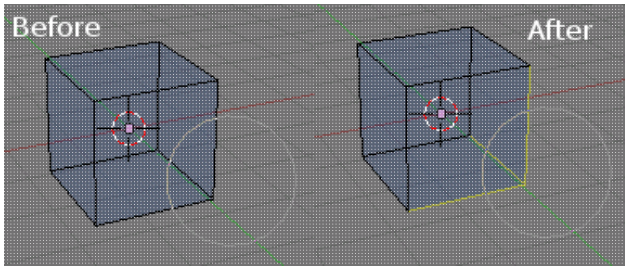
Notice in (*Complete*) that cube â← Bâ← is also selected *and* active. This means that cube â← Bâ← was the last active object prior to using the Border Select tool.

Note
Border select adds to the previous selection, so in order to select only the contents of the rectangle, deselect all with A first. In addition, you can use MMB while you draw the border to deselect all objects within the rectangle.

Circular region

This selection tool is only available in Edit mode and can be activated with B, B. That is, pressing the B key twice in a row. Once in this mode the cursor changes to a dashed cross-hair with a 2D circle surrounding it. The tool will operate on whatever the current select mode is. Clicking or dragging with the LMB, when elements are inside the circle, will cause those elements to be selected.

You can enlarge or shrink the circle region using + NumPad and - NumPad, or the Wheel.



Circle Region Select.
(*Circle Region Select*) is an example of selecting edges while in Edges select mode. As soon as an edge intersects the circle the edge becomes selected. The tool is interactive such that edges are selected while the circle region is being dragged with the LMB.

If you want to de-select elements either hold MMB or AltLMB and begin clicking or dragging again.

For Faces select mode, the circle must intersect the face indicators usually represented by small pixel squares; one at the center of each face.

To exit from this tool click RMB, or hit the Esc key.

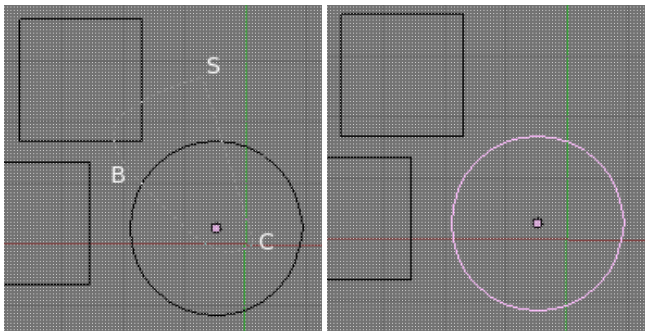
Lasso region

Lasso select is similar to Border select in that you select objects based on a region, except Lasso is a hand-drawn region that generally forms a circular/rounded shaped form; kind of like a lasso.

Lasso is available in either Edit mode or Object mode. To activate the tool use the CtrlLMB while dragging. The one difference between Lasso and Border select is that in Object mode, Lasso only selects objects where the lasso region intersects the objects centre.

To de-select use Ctrlâ § ShiftLMB while dragging.

(*Selecting*) is an example of using the Lasso select tool. Dragging started at â← Sâ← , curved around to â← Bâ← and stopped at â← Câ← . Notice that the lasso region included the circleâ← s purple coloured object centre.



Selecting. Selected.

(*Selected*) is the result with just the circle selected even though the square was in the lasso region.

Basic Mesh Modeling

In this section we will describe some of the most common mesh editing tools: Extrude, Spin, Spin Dup, Screw, Warp and To Sphere.

Each tool is described using a simple tutorial. Extrude is explained by going through a simple set of steps for making a sword. Spin is explained by making a simple wine glass. Spin Dup is explained by making the hour mark on a clock face. Screw is explained by literally making a screw. Warp is explained by warping some 3D text around a sphere. And finally, To Sphere is explained by changing a cube in a sphere.

Extrude

Mode: Edit Mode â Editing context F9

Panel: Mesh Tools â Extrude

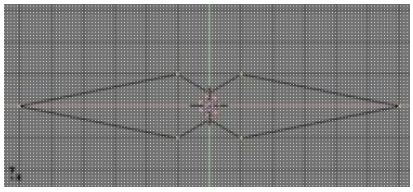
Hotkey: E

One tool of paramount importance for working with meshes is the Extrude command (E). This command allows you to create parallelepiped from rectangles and cylinders from circles, as well as easily create such things as tree limbs. Although the process is quite intuitive, the principles behind Extrude are fairly elaborate as discussed below.

- First, the algorithm determines the outside edge-loop of the extrude; that is, which among the selected edges will be changed into faces. By default, the algorithm considers edges belonging to two or more selected faces as internal, and hence not part of the loop.
- The edges in the edge-loop are then changed into faces.
- If the edges in the edge-loop belong to only one face in the complete mesh, then all of the selected faces are duplicated and linked to the newly created faces. For example, rectangles will result in parallelepipeds during this stage.
- In other cases, the selected faces are linked to the newly created faces but not duplicated. This prevents undesired faces from being retained â← insideâ← the resulting mesh. This distinction is extremely important since it ensures the construction of consistently coherent, closed volumes at all times when using Extrude.
- Edges not belonging to selected faces, which form an â← openâ← edge-loop, are duplicated and a new face is created between the new edge and the original one.
- Single selected vertices which do not belong to selected edges are duplicated and a new edge is created between the two.

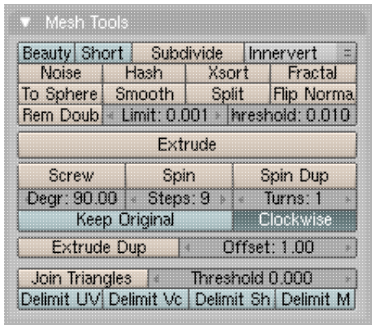
Grab mode is automatically started when the extrude algorithm terminates, so newly created faces, edges, and vertices can be moved around with the mouse. Extrude is one of the most frequently used modeling tools in Blender. Itâ← s simple, straightforward, and easy to use, yet very powerful. The following short lesson describes how to build a sword using Extrude.

The Blade

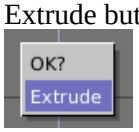


Deformed circle, to become the blade cross section.

- Start Blender and delete the default cube. In top view (7 NumPad) add a mesh circle with eight vertices. Move (G) the vertices so they match the configuration shown in (*Deformed circle, to become the blade cross section*).
- Select all the vertices (A) and scale them down (S) so the shape fits in two grid units. Switch to front view with 1 NumPad.

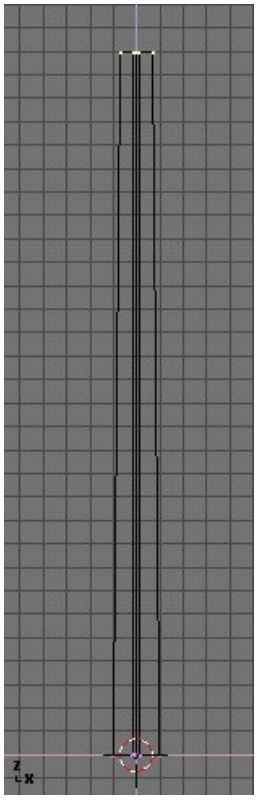


Extrude button in Editing context.



Extrude confirmation box.

- The shape weâ ← ve created is the base of the blade. Using Extrude weâ ← ll create the blade in a few simple steps. With all vertices selected press E, or click the Extrude button in the Mesh Tools panel of the Editing context (F9 â ← *Extrude button in Editing context*).
- A box will pop up asking Ok? Extrude (*Extrude confirmation box*). Click this text or press â µ Enter to confirm, otherwise move the mouse outside or hit Esc to exit. If you now move the mouse youâ ← ll see that Blender has duplicated the vertices, connected them to the original ones with edges and faces, and has entered grab mode.

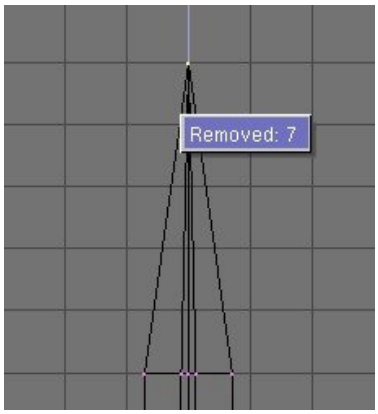


The Blade.

- Move the new vertices up 30 units, constraining the movement with Ctrl, then click LMB to confirm their new position and scale them down a little bit with S (*The Blade*).
- Press E again to extrude the tip of the blade, then move the vertices five units up. To make the blade end in one vertex, scale the top vertices down to 0.000 (hold Ctrl for this) and press W Â» Remove Doubles (*Mesh Edit Specials menu*) or click the Rem Doubles button in the Editing context (F9). Blender will inform you that it has removed seven of the eight vertices and only one vertex remains. The blade is complete (*The completed blade*)!



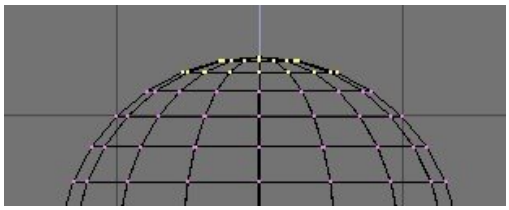
Mesh Edit Specials menu.



The completed blade.

The Handle

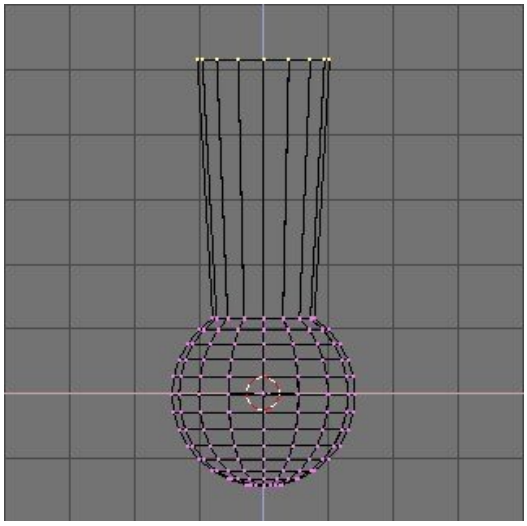
- Leave Edit mode and move the blade to the side. Add an UVsphere with 16 segments and rings and deselect all the vertices with the A.



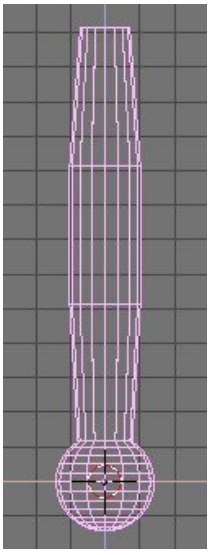
UV sphere for the handle: vertices to be removed.

- Borderselect the top three rings of vertices with B and delete them with X Â» Vertices (*UV sphere for the handle: vertices to be removed*).

- Select the top ring of vertices and extrude them. Move the ring up four units and scale them up a bit (*First extrusion for the handle*), then extrude and move four units again twice and scale the last ring down a bit (*Complete handle*).



First extrusion for the handle.

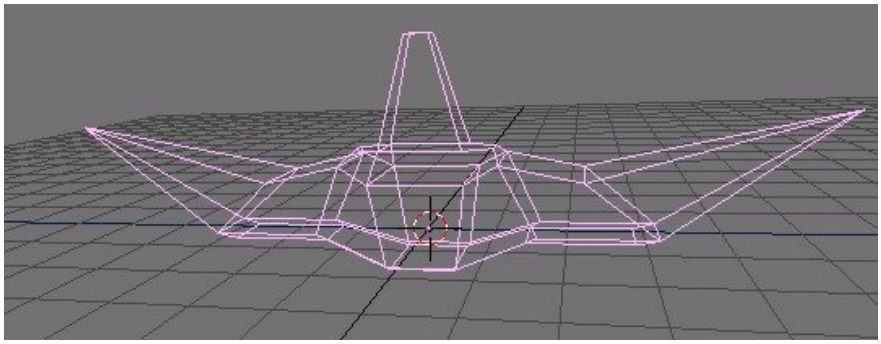


Complete handle.

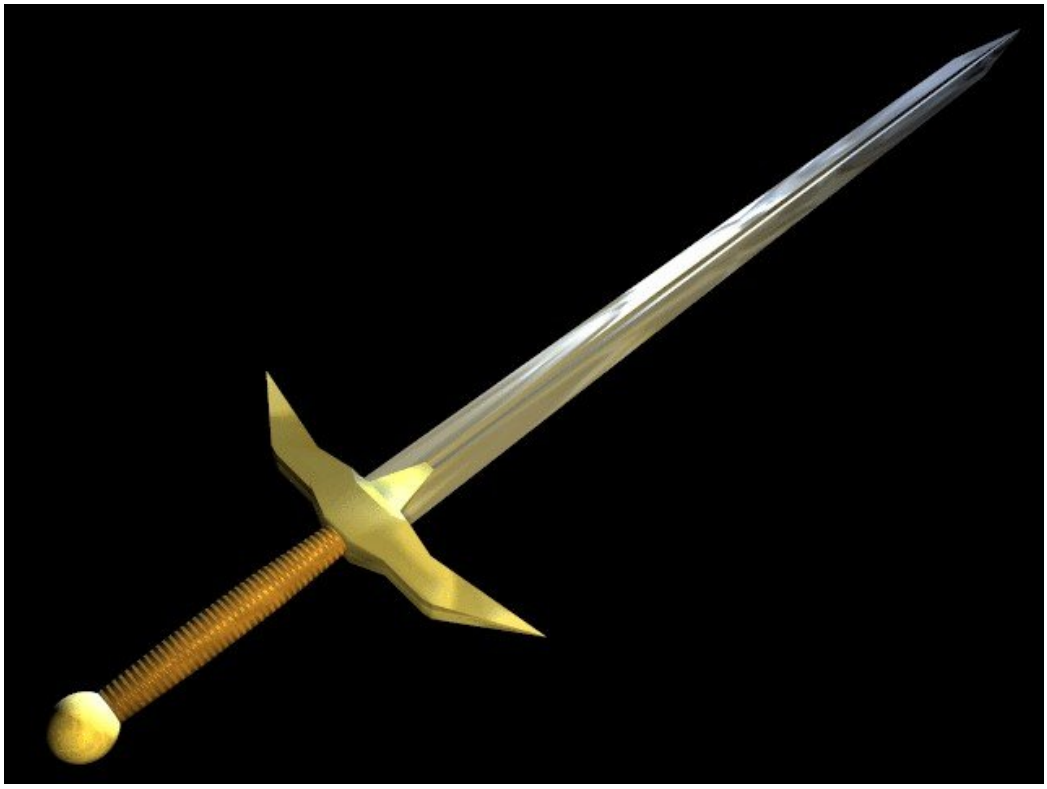
- Leave Edit mode and scale the entire handle down so that itâ← s in proportion with the blade. Place it just under the blade.

The Hilt

By now you should be used to the â← extrude â move â scaleâ← sequence, so try to model a nice hilt with it. Start out with a cube and extrude different sides a few times, scaling them where needed. You should be able to get something like that shown in (*Complete Hilt*).



Complete Hilt.
After texturing, the sword looks like (*Finished sword, with textures and materials*).



Finished sword, with textures and materials.
As you can see, Extrude is a very powerful tool that allows you to model relatively complex objects very quickly; the entire sword was created in less than half an hour. Getting the hang of â← extrude â move â scaleâ← will make your life as a Blender modeler a lot easier.

Mirror

Mode: Edit or Object modes

Hotkey: M in Edit mode; CtrlM in Object mode

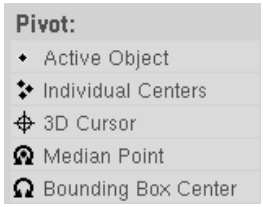
Menu: Mesh/Curve/Surface/Object Â» Mirror Â» `Axis corresponding to the wanted transformation orientation`

The mirror tool is the exact equivalent of scaling by -1 to flip Objects, Vertices, Edges, Faces around one chosen pivot point and in the direction of one chosen axis, only it is faster/handier.

Letâ← s see this in detail.

In Edit Mode

Pivot point

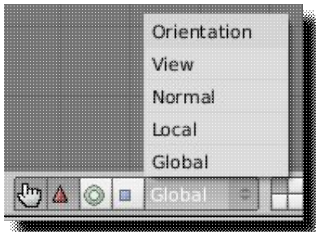


Pivot points.
Pivot points must be set first. To learn more about pivot points, see this page.

Pivot points will become the centre of symmetry.

If the widget is turned on it will always show where the pivot point is.

Transformation orientation



The available transform orientations for mirroring in Edit mode.
Transformation orientations are found on the 3D area header, next to the Widget buttons. They decide of which coordinate system will rule the mirroring.

For mirroring the available transformation orientations are:

- View, i.e. the coordinate system of the view plane of the 3D area where the transformation will occur.
- Normal, i.e. the coordinate system based on the direction and location of normals for meshes.
- Local, i.e. the coordinate system of the Object itself.
- Global, i.e. the coordinate system of the World.

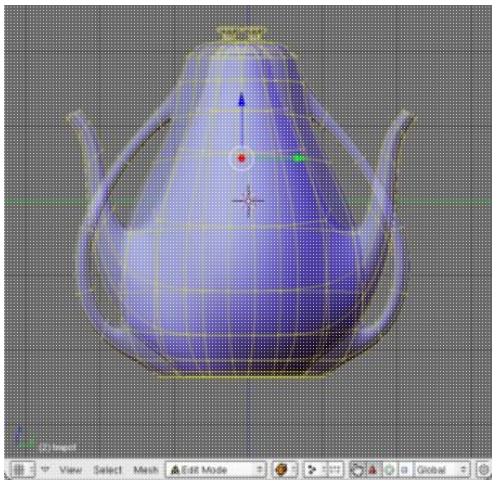
Axis of symmetry



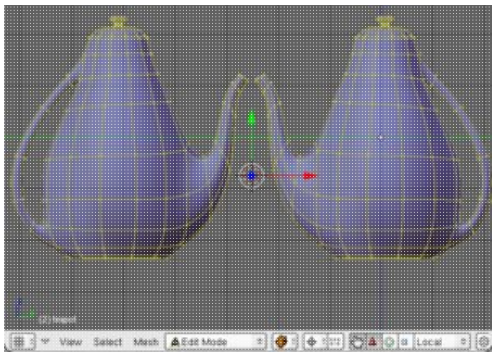
For each transformation orientation one symmetry axis.
For each transformation orientation, you can chose one of its axis along which the mirroring will occur.

As we can see the possibilities are infinite and the freedom complete: we can position the pivot point at any location around which we want the mirroring to occur, chose one transformation orientation and then one axis on it.

Here are three examples given to help figuring out what needs to be done and what result can be expected. In each case the whole geometry was duplicated with â § ShiftD and the resulting copy was mirrored.

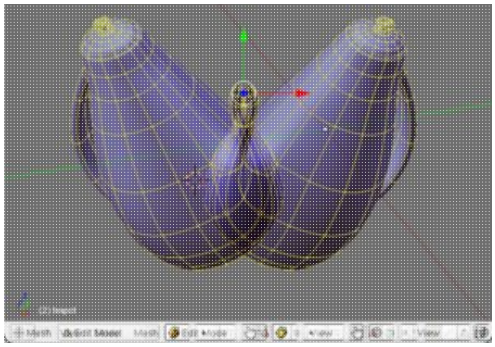


☐
Mirror around the Individual Centers and along the Global Y axis.
On (*Mirror around the Individual Centersâ←¦*) the pivot point default to **median point of the selection of vertices** in Edit mode. This is a special case of the Edit mode as explained on the pivot point page. The chosen transformation orientation is Global and the chosen axis is Y.



☐
Mirror around the 3D Cursor and along the Local X axis.

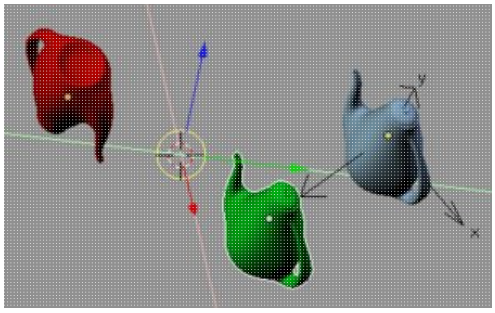
On (*Mirror around the 3D Cursorâ←¦*) the pivot point is the 3D Cursor, the transformation orientation is Local, a.k.a. the Object space, and the axis of transformation is X.



☐
Mirror around an active vertex and along the View X axis.
On (*Mirror around an active vertexâ←¦*) the vertex at the very tip of the spout is the pivot point by choosing the Active Vert/Edge/Face option and then A to select all vertices followed by RMB twice on that vertex to make it active. The transformation orientation is View and the axis of transformation is X.

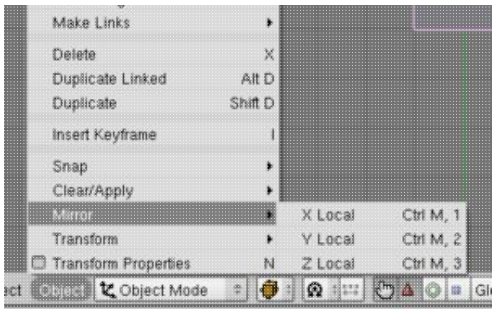
In Object Mode

Mirroring is also available in Object mode but it is limited to Local (Object space) transformations. Any other orientation gives wrong results. The following example shows what could go wrong and what a proper result should look like.

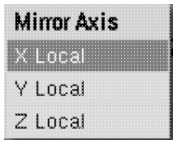


☐
Only Local space works for mirroring in Object mode.
On (*Only Local space worksâ←¦*), the red teapot is a mirrored copy of the blue one along the Global Y axis. Because the blue teapot is rotated relatively to the World this mirror resulted into an upside down copy of the original. Any transformation that is made at an angle from the Local axes of the transformed object will give wrong results.

The green teapot is also a copy of the blue one but is has been mirrored along the Local Z axis of that same blue teapot, resulting in a perfect mirror copy (colours were added afterwards).



Mirror menu in Object mode.
On (*Mirror menu in Object mode*) we can see the choice of the three Local axes which are the only ones usable. Notice also the shortcuts to the extreme right.



Pop-up menu after using CtrlM.
On (*Pop-up menu after using CtrlM*) we see the same choices which pop up in the 3D area after using the CtrlM shortcut.

In Conclusion

To summarize our survey of the Mirror tool a few recommendations:

- Do not mistake it for the Mirror modifier.
- Also remember that the results are the exact equivalent of a scale=-1 along an axis of the current transform orientation (TO). The advantage of the Mirror tool over this is that it is faster to use and almost foolproof.

Spin and SpinDup

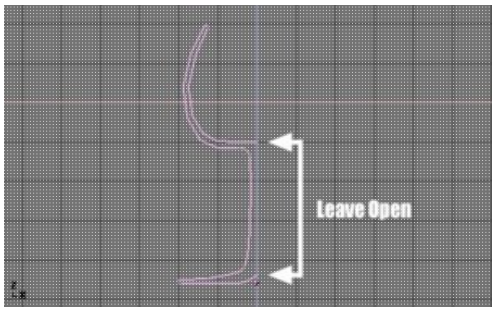
Spin and Spin Dup are two very powerful modeling tools allowing you to easily create bodies of revolution or axially periodic structures.

Spin

Mode: Edit mode â Editing context F9

Panel: Mesh Tools â Spin

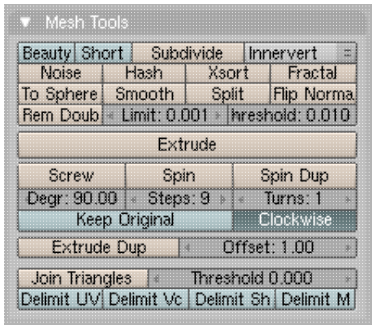
Use the Spin tool to create the sort of objects that you would produce on a lathe (this tool is often called a â← latheâ← -tool or a â← sweepâ← -tool in the literature, for this reason).



Glass profile.

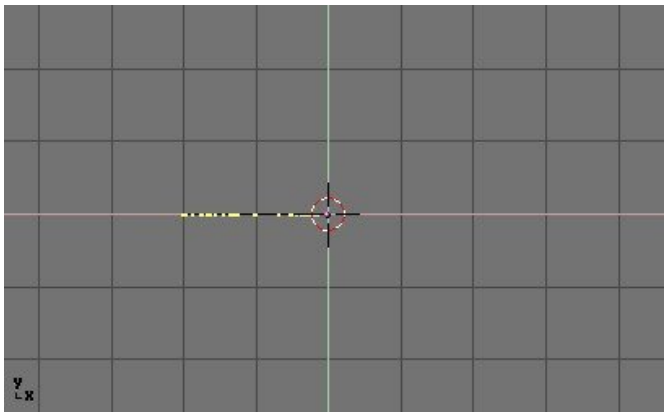
First, create a mesh representing the profile of your object. If you are modeling a hollow object, it is a good idea to thicken the outline. (*Glass profile*) shows the profile for a wine glass we will model as a demonstration.

In Edit mode, with all the vertices selected, access the Editing context (F9). The Degr button in the Mesh Tools panel indicates the number of degrees to spin the object (in this case we want a full 360Â° sweep).



Spin Buttons.
The Steps button specifies how many profiles there will be in the sweep (*Spin Buttons*).

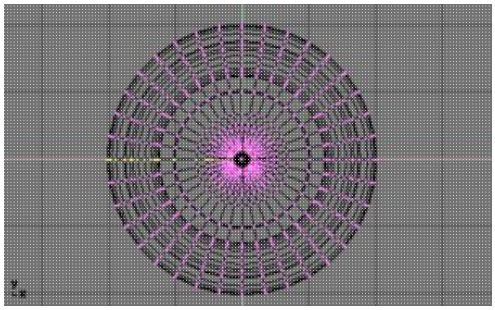
Like Spin Dup (discussed in the next section), the effects of Spin depend on the placement of the 3D cursor and which window (view) is active. We will be rotating the object around the cursor in the top view. Switch to the top view with 7 NumPad.



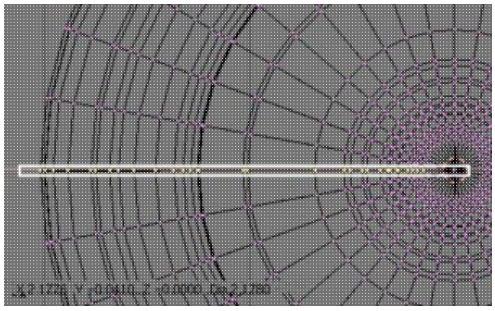
Glass profile, top view in Edit mode, just before spinning.
Place the cursor along the centre of the profile by selecting one of the vertices along the centre, and snapping the 3D cursor to that location with â § ShiftS Â» Cursor -> Selection. (*Glass profile, top view in Edit mode, just before spinning*) shows the wine glass profile from top view, with the cursor correctly positioned.



Mesh data â← Vertex and face numbers.
Before continuing, note the number of vertices in the profile. Youâ← ll find this information in the Info bar at the top of the Blender interface (*Mesh data â← Vertex and face numbers*).



Spun profile.
Click the Spin button. If you have more than one 3D view open, the cursor will change to an arrow with a question mark and you will have to click in the window containing the top view before continuing. If you have only one 3D view open, the spin will happen immediately. (*Spun profile*) shows the result of a successful spin.



Seam vertex selection.
The spin operation leaves duplicate vertices along the profile. You can select all vertices at the seam with Box select (B) shown in (*Seam vertex selection*) and perform a Remove Doubles operation.

Notice the selected vertex count before and after the Remove Doubles operation (*Vertex count after removing doubles*). If all goes well, the final vertex count (38 in this example) should match the number of the original profile noted in (*Mesh data â← Vertex and face numbers*). If not, some vertices were missed and you will need to weld them manually. Or, worse, too many vertices will have been merged.



Vertex count after removing doubles.

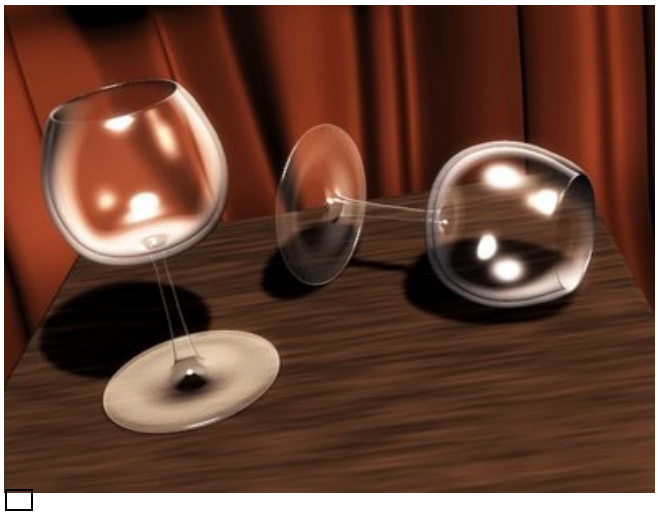


Specials menu.

Merging two vertices in one

To merge (weld) two vertices together, select both of them by â § ShiftRMB clicking on them. Press S to start scaling and hold down Ctrl while scaling to scale the points down to 0 units in the X, Y and Z axis. LMB to complete the scaling operation and click the Remove Doubles button in the Buttons window, Editing context (also available with W Â» Remove Doubles). Alternatively, you can use W Â» Merge from the same Specials menu. Then, in the new popped-up menu, choose whether the merged vertex will have to be at the center of the selected vertices or at the 3D cursor. The first choice is better in our case!

All that remains now is to recalculate the normals by selecting all vertices and pressing CtrlN and selecting Recalc Normals Outside from the pop-up menu. At this point you can leave Edit mode and apply materials or smoothing, set up some lights, a camera and make a rendering. (*Final render of the glasses*) shows our wine glass in a finished state.

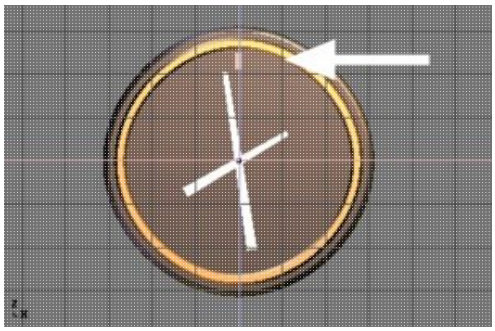


Final render of the glasses.

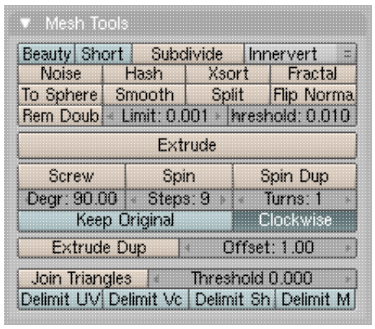
Spin Duplicate

Mode: Edit mode â Editing context F9

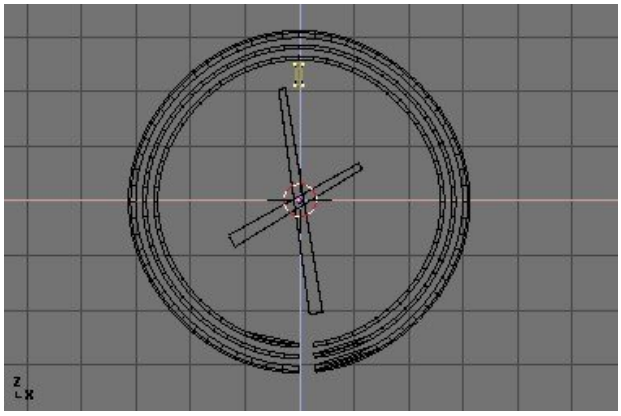
Panel: Mesh Tools â Spin Dup



Hour mark indicated by the arrow.
The Spin Dup tool is a great way to quickly make a series of copies of an object along a circle. For example, you have modeled a clock, and you now want to add hour marks. Model just one mark, in the 12 oâ← clock position (*Hour mark indicated by the arrow*). Select the mark and switch to the Editing context with F9.



Spin Dup buttons.
Set the number of degrees in the Degr: NumButton in the Mesh Tools panel to 360. We want to make 12 copies of our object, so set the Steps to 12 (*Spin Dup buttons*).

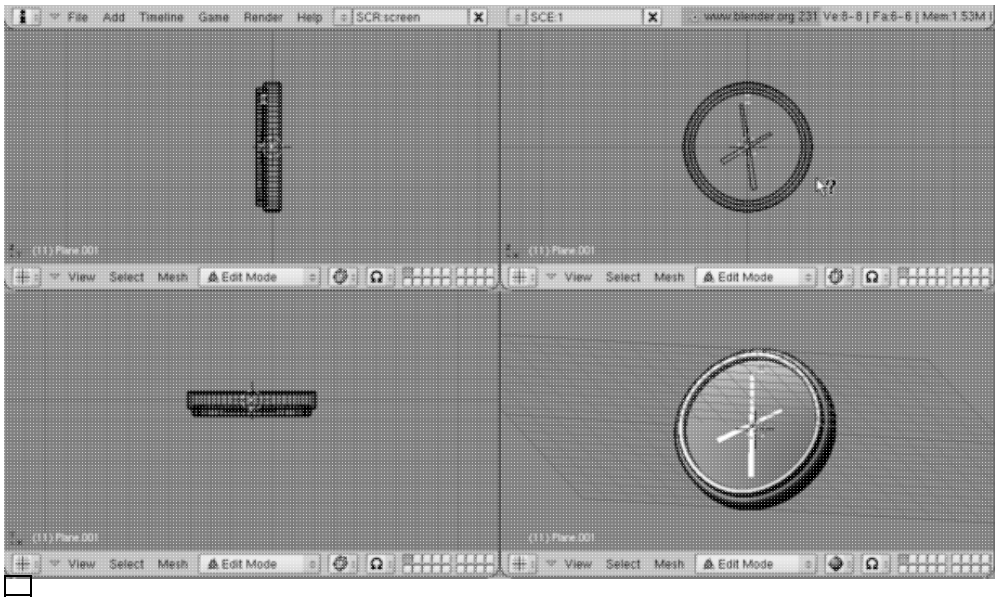


Mesh selected and ready to be SpinDuped.
Now:

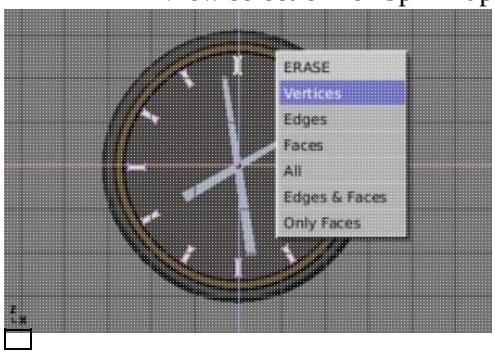
- Switch the view to the one in which you wish to rotate the object by using the keypad. Note that the result of the Spin Dup command depends on the view you are using when you press the button.
- Position the 3D cursor at the centre of rotation. The objects will be rotated around this point. Note: To place the cursor at the precise location of an existing object or vertex, select the object or vertex, and press â § ShiftS Â» Cursor -> Selection.
- Select the object you wish to duplicate and enter Edit mode with â Tab.

- In Edit mode, select the vertices you want to duplicate (note that you can select all vertices with A or all of the vertices linked to the point under the cursor with L). See (*Mesh selected and ready to be SpinDuped*).
- Press the Spin Dup button. If you have more than one 3D Window open, you will notice the mouse cursor change to an arrow with a question mark. Click in the window in which you want to perform your rotation. In this case, we want to use the front window (*View selection for Spin Dup*).

If the view you want is not visible, you can dismiss the arrow/question mark with Esc until you can switch a window to the appropriate view with the keypad.



View selection for Spin Dup.



Removal of duplicated object.

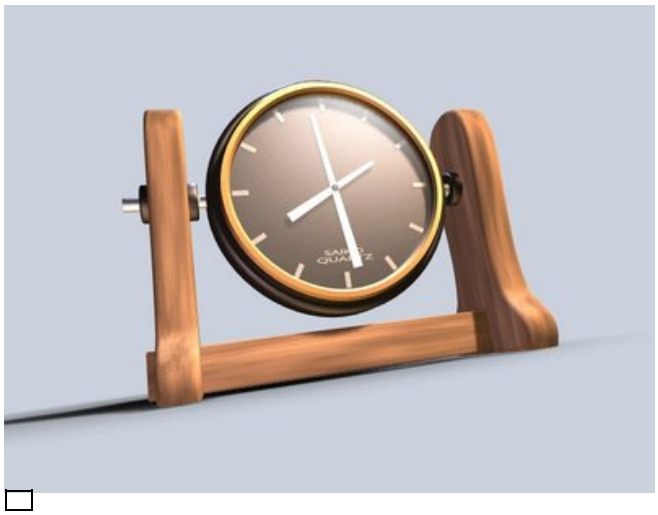
When spin-duplicating an object 360 degrees, a duplicate object is placed at the same location of the first object, producing duplicate geometry.

You will notice that after clicking the Spin Dup button, the original geometry remains selected. To delete it, simply press X Â» Vertices. The source object is deleted, but the duplicated version beneath it remains (*Removal of duplicated object*).

If you like a little math you neednâ← t bother with duplicates because you can avoid them at the start. Just make 11 duplicates, not 12, and not around the whole 360Â°, but just through 330Â° (that is `360Ã 11/12`). This way no duplicate is placed over the original object.

In general, to make `n` duplicates over 360 degrees without overlapping, just spin `n-1` objects over `360Ã (n-1)/n` degrees.

(*Final Clock Render*) shows the final rendering of the clock.



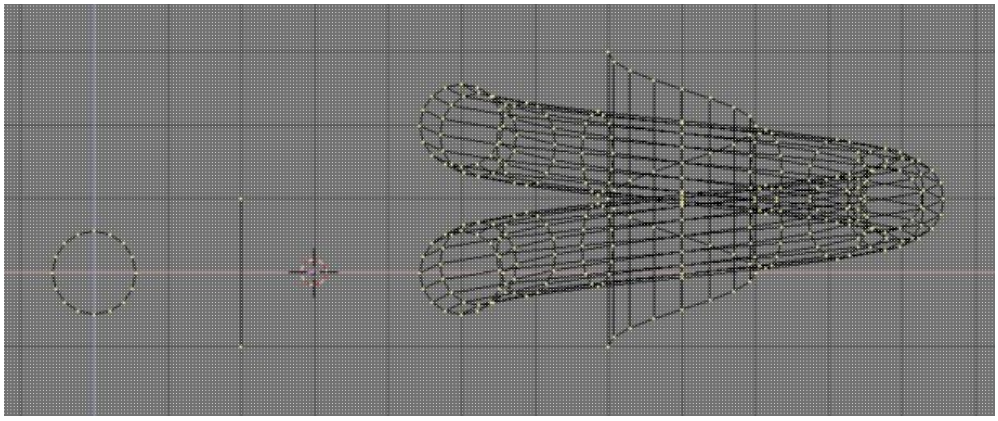
Final Clock Render.

Screw

Mode: Edit mode â Editing context F9

Panel: Mesh Tools â Screw

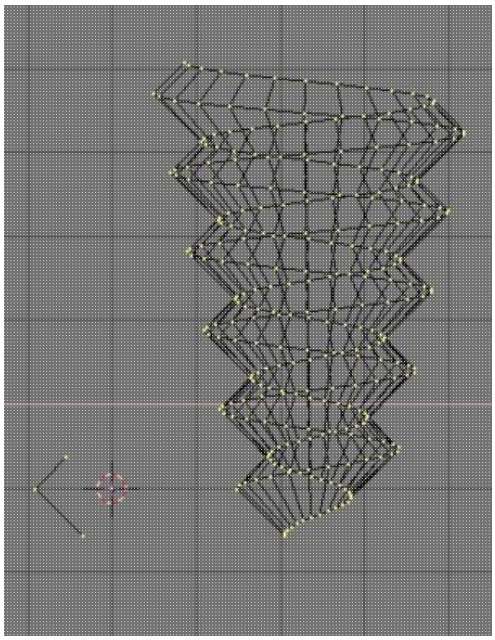
The Screw tool combines a repetitive Spin with a translation, to generate a screw-like, or spiral-shaped, object. Use this tool to create screws, springs, or shell-shaped structures.



How to make a spring: before (left) and after (right) the Screw tool.
The method for using the Screw function is strict:

- Set the 3D window to *front view* (1 NumPad).
- Place the 3D cursor at the position through which the rotation axis must pass. The rotation axis will be vertical.
- Your mesh object must contain *both* the profile to be spun and an *open line of vertices* to define how the profile is translated as it is spun. In the simplest case, the open line also serves as the profile to be spun; alternatively, a separate closed line (e.g., a circle as shown in the figure) can be specified as the profile. The open line can be a single edge, as shown in the figure, or a half circle, or whatever. You need only ensure that the line has two â← freeâ← ends (at a â← freeâ← end, a vertex is connected to only one other vertex). The Screw function uses these two points to calculate the translation vector that is added to the â← Spinâ← for each full rotation (*How to make a spring: before (left) and after*

(right) the Screw tool.). If these two vertices are at the same location, this creates a normal â← Spinâ← . Otherwise, interesting things happen!
- Select all vertices that will participate in the â← Screwâ← .
- Assign the NumButtons Steps: and Turns: in the Mesh Tools panel the desired values. Steps: determines how many times the profile is repeated within each 360Â° rotation, while Turns: sets the number of complete 360Â° rotations to be performed.
- Press Screw.



Enlarging screw (right) obtained with the profile on the left.
If there are multiple 3D Views, the mouse cursor changes to a question mark. Click on the 3D window in which the Screw is to be executed.

If the two â← freeâ← ends are aligned vertically, the result is as seen above. If they are not, the vertical component of the translation vector remains equal to the vertical component of the vector joining the two â← freeâ← vertices, while the horizontal component generates an enlargement (or reduction) of the screw as shown in (*Enlarging screw (right) obtained with the profile on the left*). In this example the open line serves as the profile as well as defining the translation.

Warp Tool

Mode: Edit mode â Editing context F9

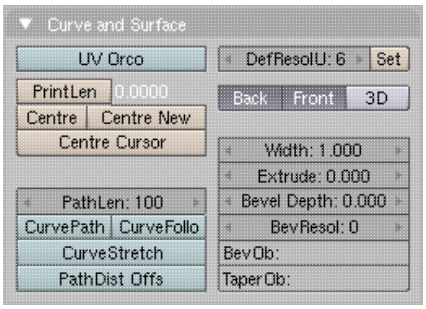
Panel: Mesh Tools â Warp

Hotkey: â § ShiftW

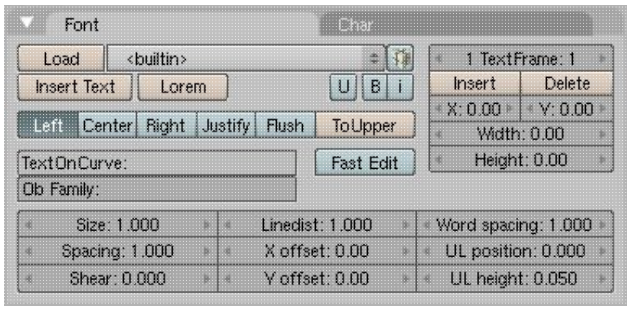
The Warp tool is a little-known tool in Blender, partly because it is not found in the Editing context of the Buttons window, and partly because it is only useful in very specific cases. At any rate, it is not something that the average Blender-user needs to use every day.

A piece of text wrapped into a ring shape is useful when creating flying logos, but it would be difficult to model without the use of the warp tool. For our example, weâ← ll warp the phrase â← Amazingly Warped Textâ← around a sphere.

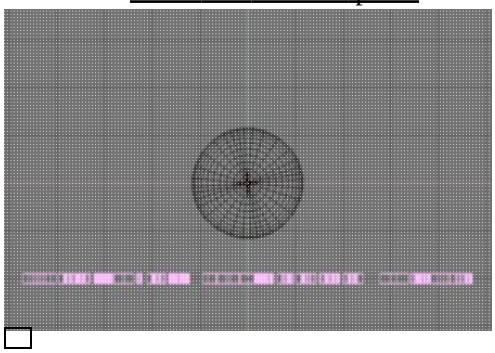
- First add the sphere.
- Then add the text in front view, in the Editing context and Curve and Surface panel set Extrude to 0.1 â← making the text 3D, and set Bevel Depth to 0.01, adding a nice bevel to the edge. Make the Bev Resol 1 or 2 to have a smooth bevel and lower the resolution so that the vertex count will not be too high when you subdivide the object later on using (*Curve and Surface panel*) and (*Font panel*).
- Convert the object to curves, then to a mesh (AltC twice), because the warp tool does not work on text or on curves.
- Subdivide the mesh twice (W Â» Subdivide Multi Â» **2**), so that the geometry will change shape cleanly, without artifacts.



Curve and Surface panel.

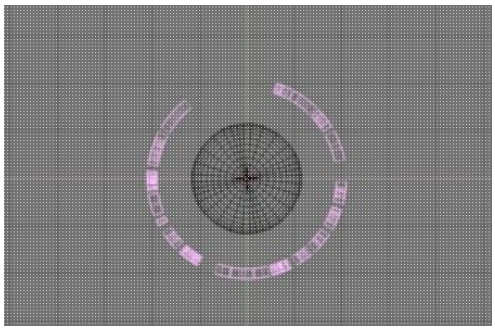


Font panel.



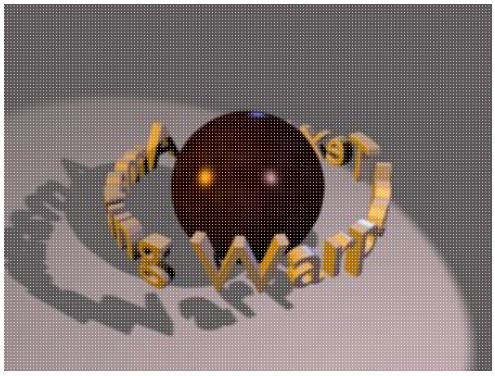
Top view of text and sphere.

Switch to top view and move the mesh away from the 3D cursor. This distance defines the radius of the warp. See (*Top view of text and sphere*).



Warped text.

Place the mesh in Edit mode (â Tab) and press A to select all vertices. Press â § ShiftW to activate the warp tool. Move the mouse left or right to interactively define the amount of warp (*Warped text*). Holding down Ctrl makes warp change in steps of five degrees.



Final rendering.
Now you can switch to camera view, add materials, lights and render (*Final rendering*).

To Sphere

Mode: Edit mode â Editing context F9

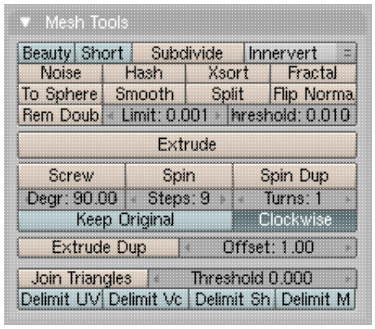
Panel: Mesh Tools â To Sphere

Hotkey: Ctrlâ § ShiftS

Another of the lesser known tools is To Sphere (Ctrlâ § ShiftS). This command allows the creation of spheres from subdivided cubes.

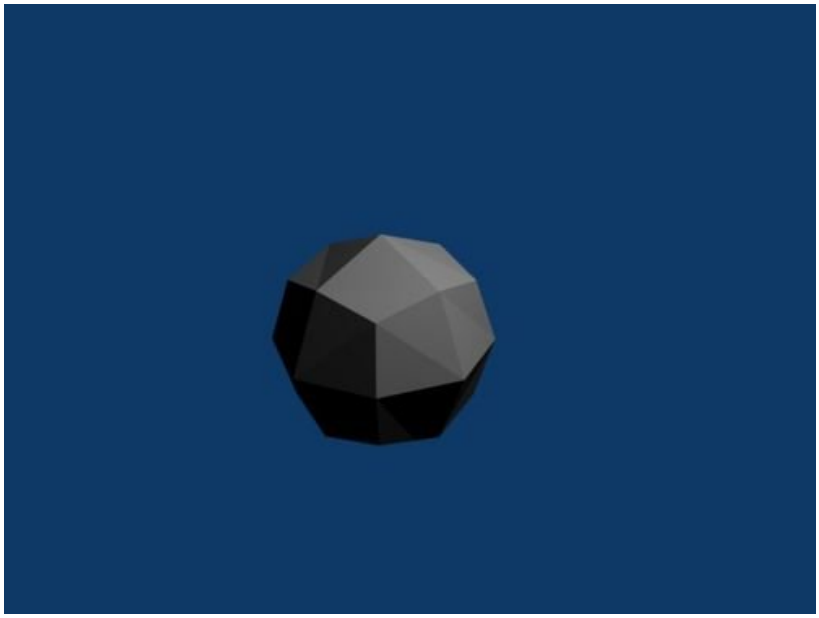
First, start with a Cube.

- Press â Tab to switch into Edit mode.



Subdivide button in Editing context.

- Make sure all the vertices of the cube are selected by pressing A twice. Then, go to the Editing context by pressing F9. You should be able to see the Mesh Tools panel now.
- Subdivide the cube by pressing the Subdivide button in the Mesh Tools panel, or with W Â» Subdivide. You can do this as many time as you want; the more you subdivide, the smoother your sphere will be.
- Click the To Sphere button now in the Mesh Tools. Select â← 100â← to make your sphere. Alternatively, you can press Ctrlâ § ShiftS and move your mouse left or right to interactively control the proportion of â← spherificationâ← (or directly type a value, like â← 1.000â← to achieve the same effect as below).

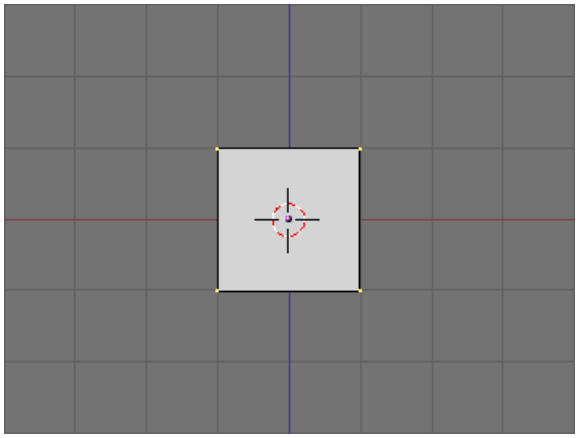


Finished low-res sphere!

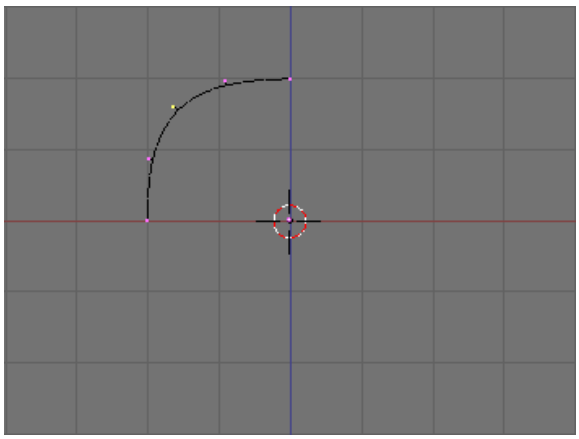
Advanced Mesh Modeling

Editorâ← s Note
This is more of a tutorial than a user manual, but we donâ← t have anything else at the moment that covers the topic. --Roger 02:17, 29 May 2007 (CEST)

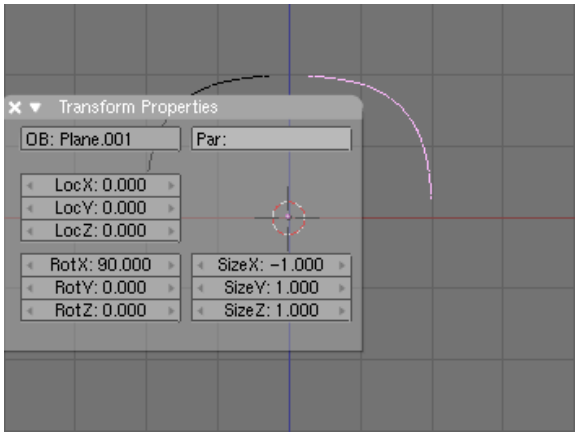
Symmetrical Modeling



A plane.
You often need to model objects which exhibit some sort of symmetry. For radial, rotational or multiple symmetry the best approach is to carefully model one base structure and then, as a last step, duplicate the base cell via Spin Dup or whichever command is most appropriate. For objects with bilateral symmetry, those with one plane of symmetry, such as most animals (humans included) and many machines, the above method implies modeling one half of the object, and then mirroring a duplicate of the first half to get the whole object. Since it is usually difficult to attain correct proportions by only modeling a half, it is possible to duplicate the half before it is completely modeled, and act on one half and automatically update the other.

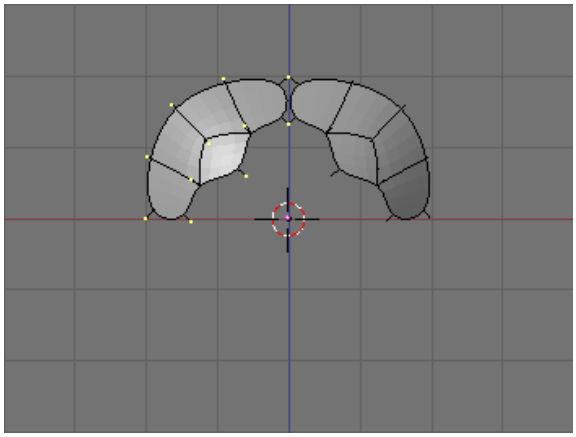


Right half.
In Front View, add a plane or whatever (*A plane*). Consider it as a starting point for one half of the object. Letâ← s say the objectâ← s right half, which for us in frontal view is on the left of the screen. The plane of symmetry is the YZ plane. Move the mesh, in Edit mode, so that it is completely on the left of centre. Delete some nodes, and add some others, to give it its general shape, as in (*Right half*).



Mirroring the linked duplicate.

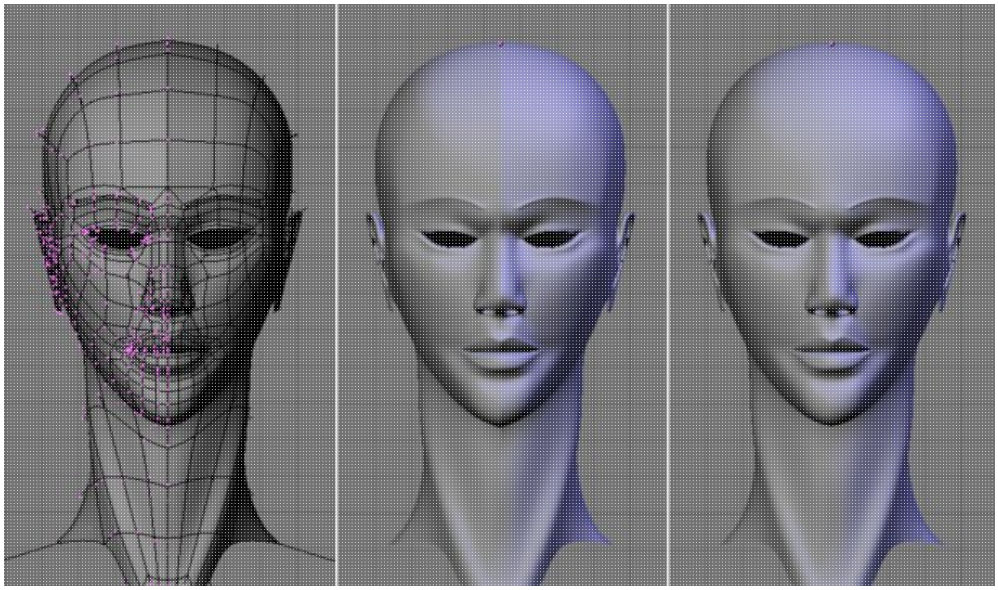
Now switch to Object mode and, with the half selected, make a *linked* duplicate with AltD. Press Esc to exit from Grab mode and press N. In the Transform Properties panel which appears, set SizeX to **-1.0** (*Mirroring the linked duplicate*). This effectively mirrors the linked duplicate with respect to the objectâ← s centre, hence the importance of keeping the centre on the plane of symmetry.



Editing one half.

Having linked-duplicated the object implies that the two objects share the *same* mesh data, which is implicitly mirrored by the unitary negative scaling along the X axis, which is normal to the symmetry plane. Now you can edit either of the two halves. Since they share mesh data, any change, be it an extrude, delete, face loop cut, etc, immediately reflects on the other side (*Editing one half*).

By carefully editing one half, possibly using a blueprint as a background to provide guidelines, very interesting results can be achieved.

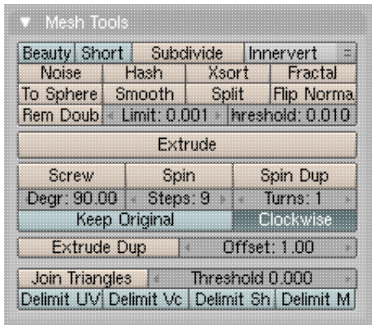


A head. Left: Edit mode; Centre: Object mode; Right: Joined.
As a final step, when symmetrical modeling is complete, the two halves must be selected and joined into a single object (CtrlJ). This makes the seam disappear, as shown in (*A head. Left: Edit mode; Center: Object mode; Right: Joined.*, centre). Once you have a single object, you can start modeling the subtle asymmetries that every natural being has.

Note
In Blender 2.33 and earlier versions the OpenGL implementation causes mirrored linked duplicates to have wrong normals, so that one of the two halves is black. This is fixed in Blender 2.34, but older versions can still use this technique by deactivating Double Sided for the mesh, while symmetrical modeling is used.

Noise

Use modifier!
Noise is quite an old feature. The recent Blender versions have a much more flexible tool to realize these sort of effects: the Displace modifier. You are strongly encouraged to use it rather than the Noise tool â← some of the key advantages of the modifier are that he can be canceled at any moment, you can precisely control how much and in which direction the displacement is applied, and much moreâ←¦



Noise button in Editing context.
The Noise function allows you to displace vertices in a mesh based on the grey-values of a texture applied to it. So, you must have a texture assigned to the material, even if that texture is not Mapped To anything. In

your texture, you should enable No RGB to convert colour textures to a gradient. You should also have subdivided your object enough to have many vertices to act on.

Use Noise to generate great landscapes or make mesh surfaces more real-world (pitted, un-smooth). The Noise function displaces vertices along the objectâ← s Â±Z-Axis only. To deform your meshâ← s other dimensions, simply rotate your object and *apply rotation*, or rotate the vertices in Edit mode, and apply Noise. Then, rotate it back again to get your original orientation.

Noise permanently modifies your mesh according to the material texture. Each click adds onto the current mesh. For a temporary effect, map the texture to Displacement for a render-time effect. In Object/Edit mode, your object will appear normal, but will render deformed.

Example

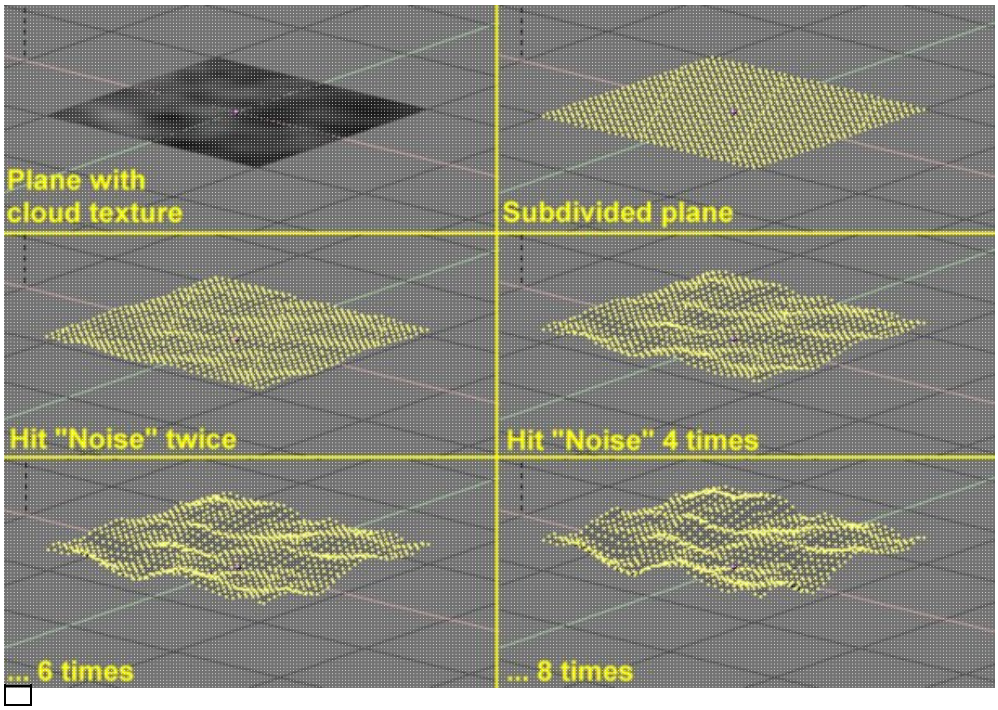


Subdivide tools in Specials menu.

Add a plane and subdivide it at least five times. To do that you can either use the Subdivide or Subdivide Multi entry in the Specials menu accessed via W; see (*Subdivide tools in Specials menu*). Using Subdivide Multi is faster and easier. Select Subdivide Multi and enter **5** for the Number of Cuts popup dialog.

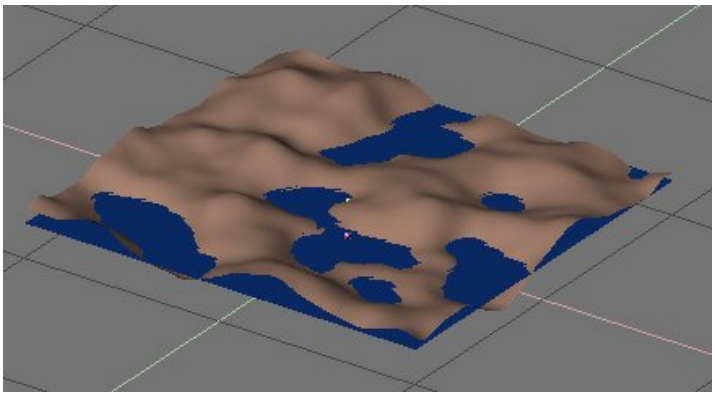
Now add a material and assign a Clouds texture to it. Adjust the NoiseSize to **0.5**. Choose white as the colour for the material and black as the texture colour, to give us good contrast for the noise operation.

Ensure that you are in Edit mode and that all vertices are selected, then switch to the Editing context (F9). Press the Noise button in the Mesh Tools panel (*Noise button in Editing context*) several times until the landscape looks nice. (*Noise application process*) is an example of applying the noise tool, showing the original â← textured â← plane as well as what happens as you press Noise. From top left to bottom right: plane with texture, sub-divided plane, Noise button hit 2, 4, 6 and 8 times.



Noise application process.

Remove the texture from the landscape now because it will disturb the look. Then add some lights, some water, smooth the terrain, and so on (*Noise generated landscape*).



Noise generated landscape.

Note

The noise displacement always occurs along the meshâ ← s *z* coordinate, which is along the direction of the *z* axis of the Object local reference.

Booleans Operations

Tip: This document section describes the *two* ways to perform Boolean operations in Blender: The â← old wayâ← (*Boolean tool*) and the â← new, improved wayâ← (*Boolean modifier*). In practice, the modifier is by far the more useful today, because of the characteristics that *all* modifiers share (and which â← old wayâ← boolean operations do not):

- Modifiers have a true user-interface. Itâ← s immediately obvious which object is being affected by the operation and which (other) object is the source.
- If you donâ← t like what you see, you can re-arrange the modifiers in the stack to change the order in which the modifiers are applied. Or you can remove or disable them, and so onâ←¦
- Like all modifiers, boolean modifiers can be selectively â← switched onâ← and â← switched offâ← .

In other words, modifiers (alone) offer *non-destructive editing*. When you use modifiers, you havenâ← t done anything *irreversible* to your model until (and unless) you choose to make the modifiers permanent by â← applyingâ← them.

Boolean Tools

Mode: Object mode (meshes only)

Panel: Modifiers (Editing context, F9)

Hotkey: W

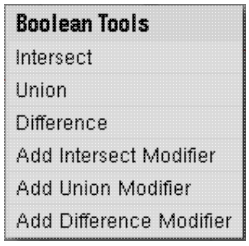
Menu: Object Â» Boolean Operation...

Description

Boolean operations are a method of combining or subtracting solid objects from each other to create a new form, a new object. Boolean operations in Blender only work on two Mesh type objects, preferably ones that are solid, or closed, with a well defined interior and exterior surface. If more than two mesh objects are selected only the *active* and previously selected object are used as operands. The boolean operations also take materials and UV-textures into account, producing objects with material indices or multi UV-mapped objects.

Options

Using the Boolean menu (W in Object mode) presents the following options:



Boolean operations.

Intersect
: Creates a new object whose surface encloses the volume *common to both* original objects.

Union
: Creates a new object whose surface encloses the *total volume of both* original objects.

Difference
: The only operation in which the order of selection is important, the active object is subtracted from the selected object. That is, the resulting object surface encloses a volume which is the volume belonging to the selected *and inactive* object, but *not* to the *active* one.

Add Intersect Modifier
: A shortcut that applies a Boolean modifier and selects Intersect in one step.

Add Union Modifier
: A shortcut that applies a Boolean modifier and selects Union in one step.

Add Difference Modifier

Guess what? A shortcut that applies a Boolean modifier and selects Difference in one step.

Boolean Modifier

Mode: Any mode

Panel: Modifiers (Editing context, F9)

Description

This sub-panel appears in the Editing context panel group which is accessed using F9 or clicking button in the Buttons window. It is part of the Modifiers parent panel. For further information about the common panel components see the Interface section on modifiers.

The downside of using the direct boolean commands is that in order to change the intersection, or even apply a different operation, you need to remove the new object and redo the command. In contrast, one can use a Boolean modifier for great flexibility and *non-destructive editing*. As with all modifiers, the booleans can be enabled/disabled or even rearranged in the stack. In addition, you can move the operands and see the boolean operation applied interactively in real time!

Caution
If the objectsâ ← meshes are too complex you may be waiting a while as the system catches up with all the mouse movements. Turning off display in the 3D View in the modifier panel can improve performance.

To get the final, â ← definitiveâ ← object from this modifier (like with the direct boolean tools) you need to â ← applyâ ← the modifier using the modifierâ ← s Apply button, and to see the result you need to move the remaining operand away or switch to local view / NumPad. Until you apply the modifier, the objectâ ← s mesh will not be modified. When you apply the boolean modifier you are notified that any mesh sticky information, animation keys and vertex information will be deleted.

Warning
There is an important difference between using boolean tools and applying a Boolean modifier: the first one **creates a new object**, whereas the second **modifies its underlying objectâ ← s mesh**. This means that when you apply a boolean modifier, **you â ← loseâ ← one of your operands**! Make a duplicate (â § ShiftD) of the object bearing the modifier, if you want to keep an â ← originalâ ← version.

However, you may not want to apply your Boolean modifier: it is indeed a very nice tool to perform some animated effects/deformations, as it is re-evaluated for each frame. This allow you e.g. to have a sort of â ← force fieldâ ← deforming around an object when this one penetrates it (see example below). Note that you can set objects to be invisible on render, which might be very useful with one of the operands of a boolean operation!

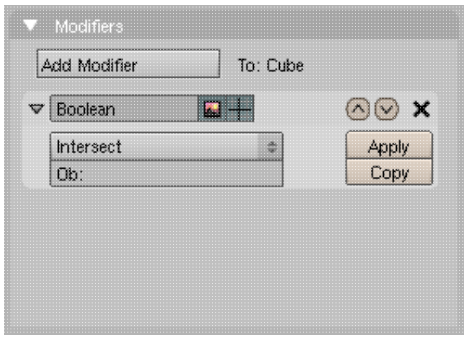
From Blender 2.49, Boolean modifiers need no more to be at the top of the modifiers stack. This was quite a heavy limitation, as it meant that you couldnâ ← t use more than one Boolean modifier at a time on a same mesh, and that you could only work on â ← rawâ ← mesh data (i.e. couldnâ ← t use e.g. the result of a Subsurf modifier as a boolean operand â ← unless you applied it first, of course).

Now, you can place a Boolean modifier where you want in the stack. You can also have multiple Boolean modifiers on the same object. And last but not least, the â ← second operandâ ← object (the one you enter the name in the Ob field of the modifier) will also be used â ← all-modifiers-appliedâ ← (e.g. you can use its subsurfed form).

Warning

All these new possibilities are great. But be careful: the boolean operations are still a complex process, which may take a very long time to complete on operands with many vertices. So take attention to all modifiers placed *before* the Boolean one, especially to Subsurf ones!

Options



Modifiers panel with Boolean modifier activated.

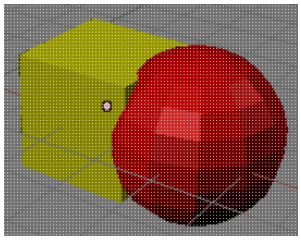
Operation
: This drop-down lists available boolean operation types (Intersect, Union or Difference).

Ob
: The name of the object to be used as the second operand to this modifier.

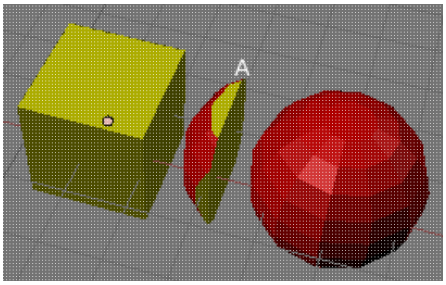
Examples

Intersect

The cube and the sphere have been moved to reveal the newly created object (â← Aâ←). Each face of the new object has the material properties of the corresponding surface that contributed to the new volume based on the Intersect operation.



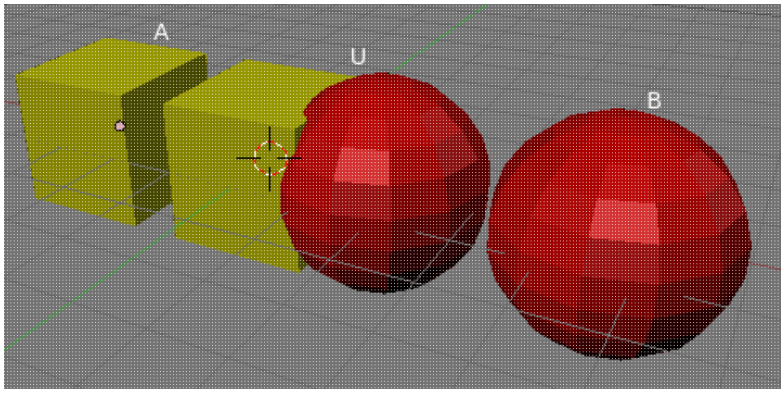
Before Intersect.



After Intersect.

Union

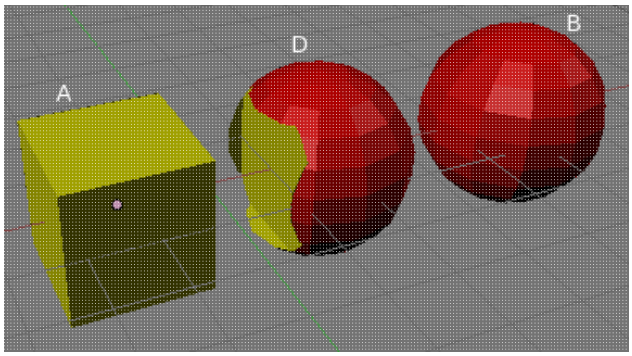
The cube (â← Aâ←) and the sphere (â← Bâ←) have been moved to reveal the newly created object (â← Uâ←). â← Uâ← is now a single mesh object and the faces of the new object have the material properties of the corresponding surface that contributed to the new volume based on the Union operation.



Union example.

Difference

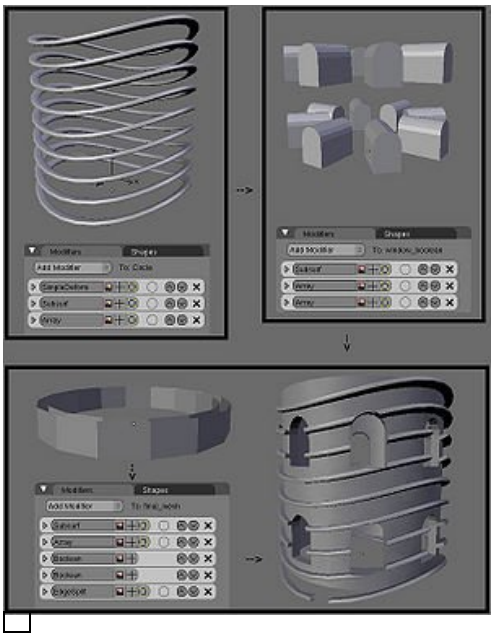
The Difference of two objects is (as the â← standardâ← subtraction) not commutative in that the *inactive* object minus the *active* object (behavior of this operation) does not produce the same as *active* minus *inactive*. The *active* cube (â← Aâ←) has been subtracted from an *inactive* sphere (â← Bâ←), and both have been moved to reveal the newly created object (â← Dâ←). â← Dâ← is now a single mesh object and the faces of the new object have the material properties of the corresponding surface that contributed to the new volume based on the Difference operation. Dâ← s volume is less than Bâ← s volume because it was decreased by subtracting part of the cubeâ← s volume.



Difference example.

More complex examples with Boolean modifiers

Multiple Modifiers



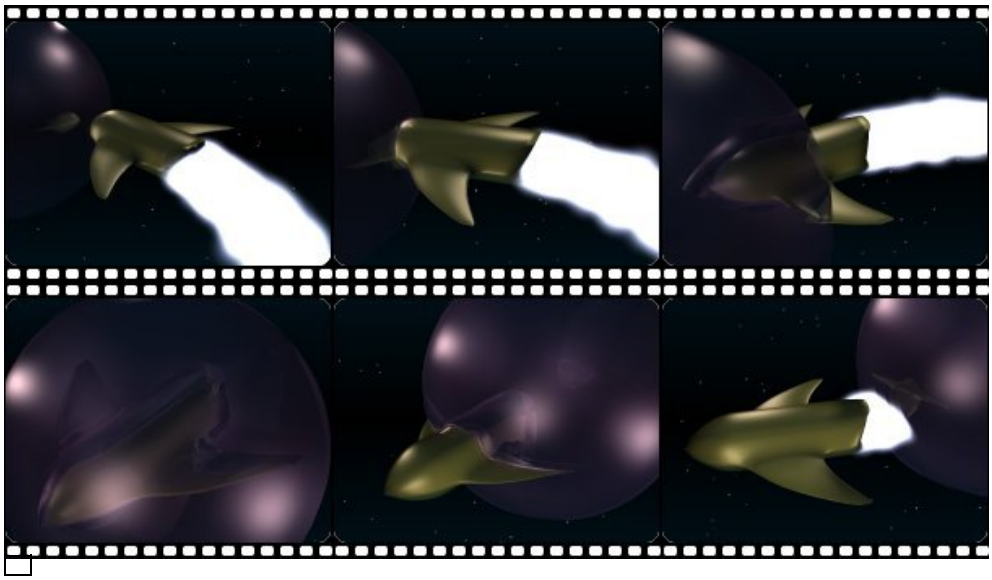

A complex object modeled using multiple Boolean modifiers.

The process to model the complex object shown to the right.
(*The process to model the complex object shown to the right*) illustrates how to model quite a complex mesh object, using only three basic forms, and several modifiers. As usual, click on the picture to enlarge.

(*A complex object modeled using multiple Boolean modifiers*) is the result of this modeling process. You can download the demo blender file here.

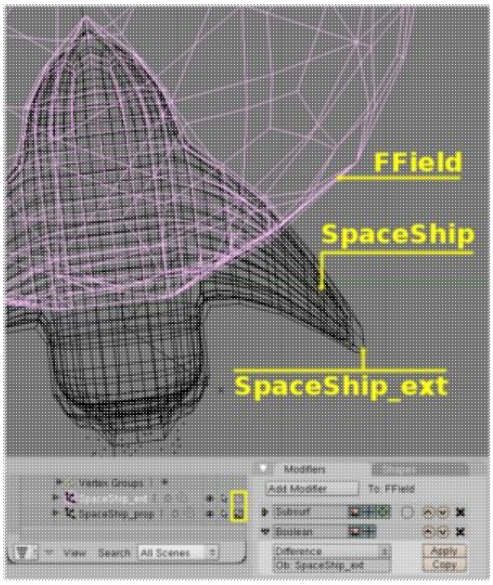
Here is a video showing the new Blender 2.49 workflow:

Animation



An animated boolean deformation.

(*An animated boolean deformation*) shows some key-frames of the boolean-deformation of a â← force fieldâ← sphere penetrated by a sort of space swallow. It tries to demonstrate the possibilities of Boolean modifiers in animations.



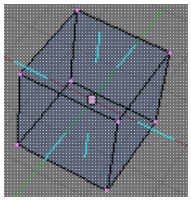
The schema of the effect.

(*The schema of the effect*) summarizes the key points of the boolean-deformation effect. There is one small trick here: one of the boolean operands (â← `SpaceShip_ext`â←) is not rendered (the third toggle control, at the right of the Outliner window, in Outliner display), which allows us to keep the â← `FField`â← surface a bit away from the real â← `SpaceShip`â← (else, we would have a lot of artifacts!).
â← `SpaceShip_ext`â← is a linked-duplicate (AltD) of the â← `SpaceShip`â← object, just slightly up-scaled and parented to the â← originalâ← .

You can download the blender file here.

And here is the complete video of this anim (Media:manual-modifiers_boolean_exvideo.avi):

Technical Details



Visible normals.
The boolean operations rely heavily on the surface normals of each object and so it is very important that the normals are defined properly and consistently. This means each objectâ← s normals should point outward. A good way to see the objectâ← s normals is to turn on the visibility of normals using the Mesh Tool More panel (Buttons window, Editing context â← Â F9) and clicking Draw normals. The normals are only visible while in Edit mode. (*Visible normals*) is an example of a cube with its normals visible.

In the case of *open* objects, that is objects with holes in the surface, the interior is defined mathematically by extending the boundary faces of the object to infinity. As such, you may find that you get unexpected results for these objects. A boolean operation never affects the original objects, the result is always a new object.

Warning
This is **not true** with the Boolean modifiers: when they are applied, they modify their owner object, and do not create a new one!

Some operations will require you to move the operands or switch to local view (/ NumPad) to see the results of the boolean operation.

Limitations & Workarounds

The number of polygons generated can be very large compared to the original meshes, especially when using complex concave objects. Furthermore, the polygons that are generated can be of poor quality, for example, very long and thin and sometimes very small. Try using the Decimate modifier (Editing context, F9) to fix this problem.

Sometimes the boolean operation can fail with a message saying â← `An internal error occurred -- sorry`â← . If this occurs, try to move or rotate the objects just a very small amount and try again.

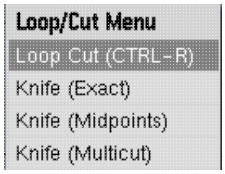
Edge Tools

Mode: Edit mode (Mesh)

Hotkey: CtrlE or K

Menu: Mesh Â» Edges

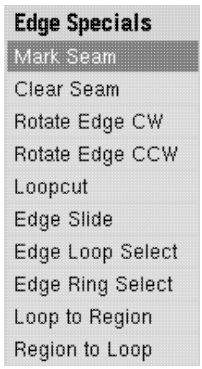
Introduction



Loop/Cut Menu.
A key issue in modeling is the necessity to add vertices in certain zones of the mesh, and this is often regarded as splitting, or adding, edges in a given region.

Loops often play an important role in this process. For a brief introduction to loops, please refer to Edge and Face Loops.

Many edge tools are grouped in a menu which is linked to K hotkey (*Loop/Cut Menu*), but each individual tool has its own hotkey as well.



Edge Specials menu.
Edge Tools for selection and manipulation are grouped in a menu which is linked to CtrlE hotkey (*Edge Specials menu*).

Edge Slide, for example, slides the vertices in the loop along the edges. If you select a loop on an egg-shaped object, for example, sliding the vertices will move them, not left/right or up/down, but instead proportionally move them as if they were sliding along the edge using the edge as a guide. More information on selecting loops is found below.

Tools donâ← t work on modifiers
In general, you cannot use any tool on a mirrored side, as that side is just a mirror image of the primary side. Tools also do not work on subsurfed or multires â← edgesâ← shown; use the tools by working on vertices/edges/faces on the primary part of your mesh when using a modifier.

Edge Selection

Mode: Edit mode (Mesh)

Hotkey: RMB

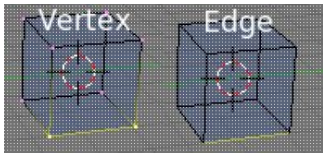
Description



Select modes.
In Edit mode, there are a few ways to select edges: implicitly, explicitly, looping or by region selection.
Implicit means you describe a more complex element using less complex elements. For example, to describe an edge you need to specify two vertices and to describe a face you need to specify three or more vertices or three or more edges.

Region selection is a tool that allows selection of edges and faces based on an intersection with a rectangular, circular or lasso 2D region.

Explicit Edge Selection
: To select an edge use edge select mode and RMB click on an edge. To select additional edges use â § ShiftRMB .



Implicit Edge selection.

Implicit Edge Selection
: The other way to select an edge is to select two vertices that bound the edge of interest. You are implying which edge by selecting its bounding vertices. To select an edge implicitly use vertex select mode in combination with RMB and/or â § ShiftRMB .
In (*Implicit Edge selection*), the cube on the left shows an edge selected using vertices. The cube on the right is what shows when you switch to edge select mode.

Options

If you are in solid, shaded, or textured viewport shading mode (not bounding box or wireframe), you will have a fourth button that looks like a cube. When enabled, this limits your ability to select based on visible edges (as if the object was solid), and prevents you from accidentally selecting, moving, deleting or otherwise working on backside or hidden edges.

Edge Loop Selection

Mode: Edit mode (Mesh) â Vertex or Edge select mode

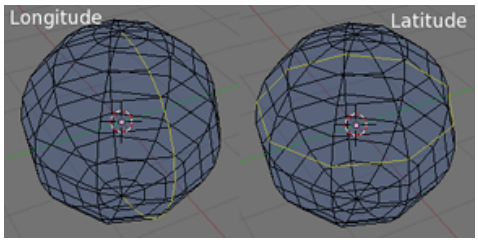
Hotkey: AltRMB or CtrlE Â» 7 (based on existing edge selection)

Menu: Select Â» Edge Loop (based on existing edge selection)

Description

Holding Alt while selecting an edge selects a loop of edges that are connected in a line end to end, passing through the edge under the mouse pointer. Holding Altâ § Shift while clicking adds to the current selection. Edge loops can also be selected based on an existing edge selection, using the Edge Specials Â» Edge Loop Select menu option (CtrlE).

Examples



Longitudinal and latitudinal edge loops.
The left sphere shows an edge that was selected longitudinally. Notice how the loop is open. This is because the algorithm hit the vertices at the poles and terminated because the edge at the pole connects to more than 3 other edges. However, the right sphere shows an edge that was selected latitudinally and has formed a closed loop. This is because the algorithm hit the first edge that it started with.

Technical Details

The algorithm for selection is as follows:

- First check to see if the selected element connects to only 3 other edges.
- If the edge in question has already been added to the list, the selection ends.
- Of the 3 edges that connect to the current edge, the ones that share a face with the current edge are eliminated and the remaining edge is added to the list and is made the current edge.

Edge Ring Selection

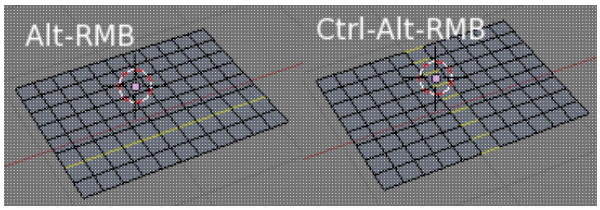
Mode: Edit mode (Mesh) â Edge select mode

Hotkey: CtrlAltRMB

Description

In edge select mode, holding CtrlAlt while selecting an edge selects a sequence of edges that are not connected, but on opposite sides to each other continuing along a face loop. Using the same command in vertex select mode will select such a face loop implicitly.

Examples



A selected edge loop, and a selected edge ring.
In (*A selected edge loop, and a selected edge ring*), the same edge was clicked on but two different â← groups of edgesâ← were selected, based on the different commands. One is based on edges during computation and the other is based on faces.

Technical Details

Edge ring selection is based on the same algorithm as in Face Loop Selection, though the end results differ as only some edges are selected.

Rotate Edge CW / Rotate Edge CCW

Mode: Edit mode (Mesh)

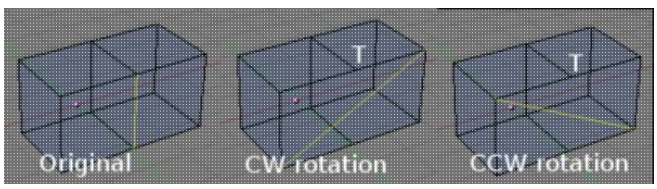
Hotkey: CtrlE Â» 3 and CtrlE Â» 4

Menu: Mesh Â» Edges Â» Rotate Edge CW / Rotate Edge CCW

Description

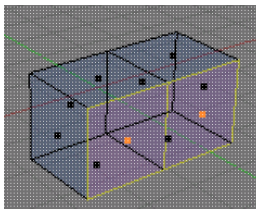
Rotating an edge clockwise or counter-clockwise spins an edge between two faces around their vertices. This is very useful for restructuring a meshâ← s topology. The tool can operate on one explicitly selected edge, or on two selected vertices or two selected faces that implicitly select an edge between them.

Examples



Selected edge rotated CW and CCW.

Be aware that sometimes, as shown in (*Selected edge rotated CW and CCW*), indicated with a â← Tâ← , that you could produce what appears to be â← Tâ← junctions/nodes by using this tool. However, Blender has created additional edges that prevent cracks in the mesh. You can test this by selecting the vertex at the â← Tâ← and moving it around while noting that there are two edges now instead of one long edge.



Adjacent selected faces.
To rotate an edge based on faces you must select two faces, (*Adjacent selected faces*), otherwise Blender notifies you with an error message, â← `ERROR: Select one edge or two adjacent faces`â← . Using either Rotate Edge CW or Rotate Edge CCW will produce exactly the same results as if you had selected the common edge shown in (*Selected edge rotated CW and CCW*.).

Edge Slide

Mode: Edit mode (Mesh) â Vertex or Edge select mode

Hotkey: CtrlE Â» 6 NumPad

Menu: Mesh Â» Edges Â» Slide Edge

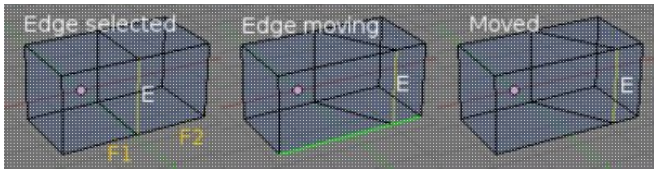
Description

Edge Slide slides one or more edges along faces adjacent to the selected edge(s) with a few restrictions involving the selection of edges.

Options

- LMB confirms the tool, and RMB or Esc cancels.
- This tool has both a *Percentage* and *Proportional* mode, which is displayed in the 3D View header. These modes behave the same as with the Loop Subdivide tool, including all keys for controlling precision edge movement.

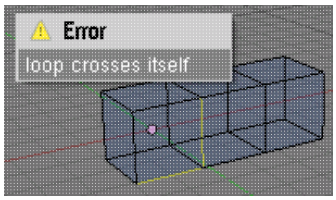
Examples



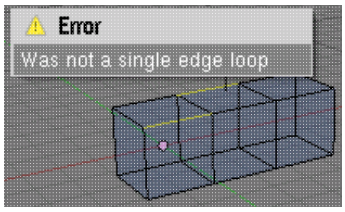
Simple edge slide.

(*Simple edge slide*) is an example of sliding an edge along an extruded box. The selected edge is labeled â← Eâ← and the adjacent faces to that edge are â← F1â← and â← F2â← . In â← Edge movingâ← , the edge is being slid along the edge drawn in green. â← Movedâ← shows the results.

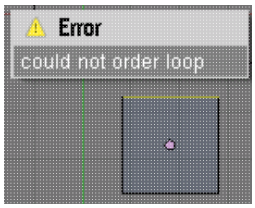
Limitations & Workarounds



Loop crosses.



Single edges.



Order loop.

There are restrictions on the type of edge selections that can be operated upon. Invalid selections are:

â← Loop crosses itselfâ←
: This means that the tool could not find any suitable faces that were adjacent to the selected edge(s). (*Loop crosses*) is an example that shows this by selecting two edges that share the same face. A face cannot be adjacent to itself.

â← Was not a single edge loopâ←
: Most likely you have selected edges that donâ← t share the same edge loop. (*Single edges*) is an example where the selected edges are not in the same edge loop, which means they donâ← t have a common edge. You can minimize this error by always selecting edges end to end or in a â← Chainâ← .

â← Could not order loopâ←
: This means the tool could not find an edge loop based on the selected edge(s). (*Order loop*) is an example where a single edge was selected in a 2D *Plane* object. An edge loop can not be found because there is only one face. Remember, edge loops are loops that span two or more faces.

A general rule of thumb is that if multiple edges are selected they should be connected end to end such that they form a continuous chain. This is *literally* a general rule because you can still select edges in a chain that are invalid because some of the edges in the chain are in different edge loops. (*Loop crosses*) is just such an example where the selected edges form a chain but they are not in the same edge loop.

If you select multiple edges just make sure they are connected. This will decrease the possibility of getting looping errors.

Bevel

Mode: Edit mode (Mesh)

Hotkey: W Â» Alt2

Menu: Mesh Â» Edges Â» Bevel

Description



With bevel and without bevel.
A bevel is something that smooths out a sharp edge or corner. True world edges are very seldom exactly sharp. Not even a knife blade edge can be considered perfectly sharp. Most edges are intentionally beveled for mechanical and practical reasons.

Bevels are also useful for giving realism to non-organic models. In the real world, the blunt edges on objects catch the light and change the shading around the edges. This gives a solid, realistic look, as opposed to un-beveled objects looking false.

Use modifiers!
As the Booleans, the Bevel tool is an heritage of old Blender versions. Nowadays, you should use its modifier version, which gives you the same options, with additional goodies, like the bevel width controlled by the vertices weight, and all the modifiers general enhancements (non-destructive operations, â←¦). Note that the Bevel modifier has no recursive option. However, this is not a problem, as you can add as much modifiers as you need â← this is in fact even more flexible!

Options

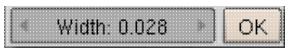


Recursion level.

Recursion
: The number of recursions in the bevel can be defined in an additional popup number field. The greater the number of recursions, the smoother the bevel.
If it is one, then each face is reduced in size and each edge becomes a single new face. *Tri* and *quad* faces are created as necessary at the corresponding vertices. If the Recursion number is greater than one, then the bevel procedure is applied that number of times. Hence, for a Recursion of **2** each edge is transformed into 4 edges, three new faces appear at the edge while smoothing the original edge. In general the number of new edges is 2 elevated to the Recursion power.

Width


: Bevel window header.
You can change the bevel width by moving the mouse towards and away from the object. As usual, the scaling can be controlled to a finer degree by either holding Ctrl, to scale in 0.1 steps, or by holding â § Shift to scale in 0.001 steps. LMB finalizes the operation, RMB or Esc aborts the action.



: Alternatively, you can manually enter a scaling value by pressing Space. A popup dialog appears, asking you to type in the beveling scale factor labeled as Width. The scale is limited to a range from

0.0 to **10.0** and upon hitting OK the bevel action is completed.

Hints

Remember that in each recursion, for each new edge two new vertices are created, with additional vertices created at the intersection between edges. This means your vertex count can quickly become enormous if you bevel with a high recursion!

Bevel Selected Edges

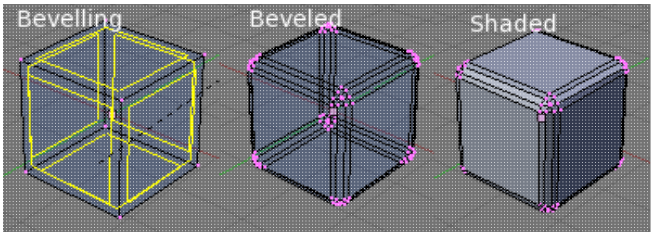
With the W Â» Bevel command, all edges of a given mesh are beveled. To bevel only selected edges, use the â← Bevel Centerâ← script instead.

As already pointed, you are encouraged to use the Bevel modifier and its bevel weight option to control which edges are beveled, and to which extent.

Examples

(*Beveling a cube*) is an example of beveling a cube with a Recursion of **2**. Once the Recursion number is set each face of the mesh receives a yellow highlight. The cube labeled â← `Bevelling`â← is the tool in action.

The final result can be seen in the cube labeled â← `Beveled`â← or â← `Shaded`â← .



Beveling a cube.

Face Tools

Face Selection

Mode: Edit mode (Mesh)

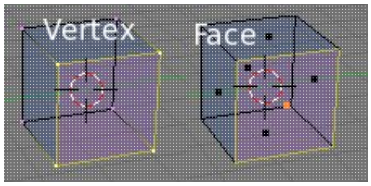
Hotkey: RMB

Description



Select modes.
In Edit mode there are a few ways to select **faces**: implicitly, explicitly, looping or by region. *Implicit* means you describe a more complex element using less complex elements. For example, to describe an edge you need to specify two vertices and to describe a face you need to specify three or more vertices or three or more edges.

Explicit Face Selection
: To select a face use the Face select mode (*Select modes*) and RMB. To select additional faces use â § ShiftRMB .



Implicit Face selection.

Implicit Face Selection
: Selecting the three or four vertices that bound the face of interest in vertex select mode selects it implicitly. (*Implicit Face selection*) shows a face selected on the left cube using its vertices. The cube on the right is what shows when you switch to face select mode.
You can also implicitly select faces by selecting the bounding edges of the face of interest. This will produce the same results as selecting vertices.

Deleting faces may delete edges and vertices
If a vertex that defines a selected face is not connected to anything else, and you delete the face, Blender deletes the vertex as well as its connecting edges. This is so you donâ ← t end up with a bunch of stray vertices running around unconnected in 3D space. If you want a vertex and/or an edge to stay, you have first to see it (vertices and edges are both visible in vertex select mode; only the edges are visible in edge select mode; neither the vertices nor the edges are visible in face select mode). Then you must use X Â» Faces Only or X Â» Edges and Faces, accordingly to what you want to keep of what you see.

Options

If you are in solid, shaded, or textured viewport shading mode (not bounding box or wireframe), you will have a fourth button that looks like a cube. When enabled, this limits your ability to select from visible faces (as if the object was solid), and prevents you from accidentally selecting, moving, deleting or otherwise working on backside or hidden faces.

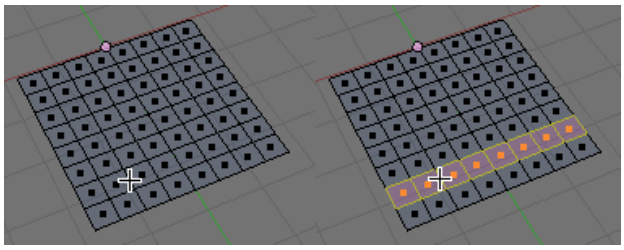
Face Loop Selection

Mode: Edit mode (Mesh) â Face or Vertex select modes

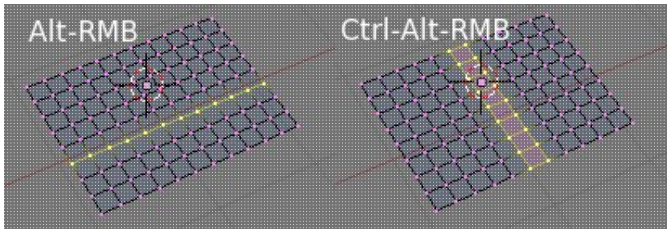
Hotkey: AltRMB (face select mode) or CtrlAltRMB (vertex select mode)

In face select mode, holding Alt while selecting an **edge** selects a loop of faces that are connected in a line end to end, along their opposite edges. In vertex select mode, the same can be accomplished by using CtrlAlt to select an edge ring, which selects the face loop implicitly.

Examples



Face loop selection.
This face loop was selected by clicking with AltRMB on an edge, in face select mode. The loop extends perpendicular from the edge that was selected.



Alt versus CtrlAlt in vertex select mode.
A face loop can also be selected in Vertex select mode, see (*Alt versus CtrlAlt in vertex select mode*). The edges selected in the grid labeled â← `Alt-RMB`â← is a result of selecting and edge loop versus selecting an edge ring. Because in vertex select mode, selecting opposite edges of a face implicitly selects the entire face, the face loop has been selected implicitly.

Note that in these cases the generated result of the algorithm was *vertices* because we were in Vertex select mode. However, had we had been in Edge select mode, the generated result would have been selected edges.

Technical Details

The algorithm for selection is as follows:

- A face loop is made by two neighbouring edge loops.
- Extends only to quadrilateral faces.
- Ends when a triangular face is met (and the two bounding edge loops merge into one).

Advanced Edge and Face Tools

Loop to Region and Region to Loop

Mode: Edit mode (Mesh)

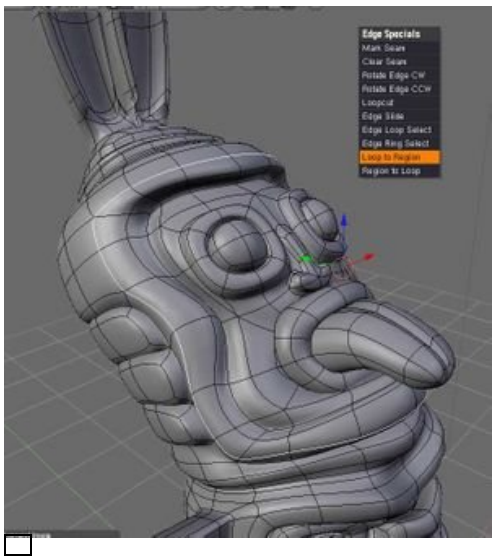
Hotkey: CtrlE Â» 9 and CtrlE Â» 0

Menu: Select Â» Loop to Region / Region to Loop

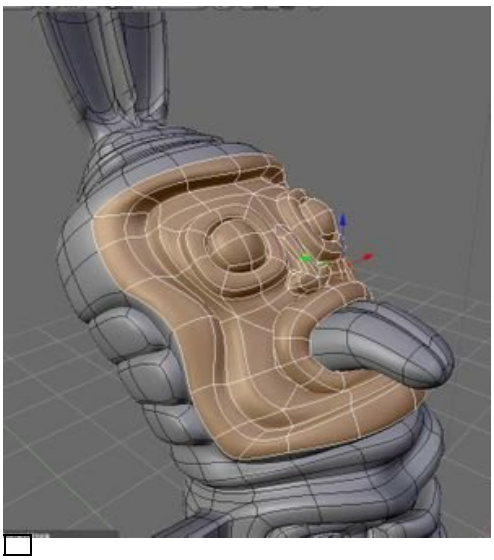
Description

This tool examines the current set of selected edges and separates them into groups of â← loopsâ← that each bisect the mesh into two parts. Then for each loop it selects the smaller â← halfâ← of the mesh.

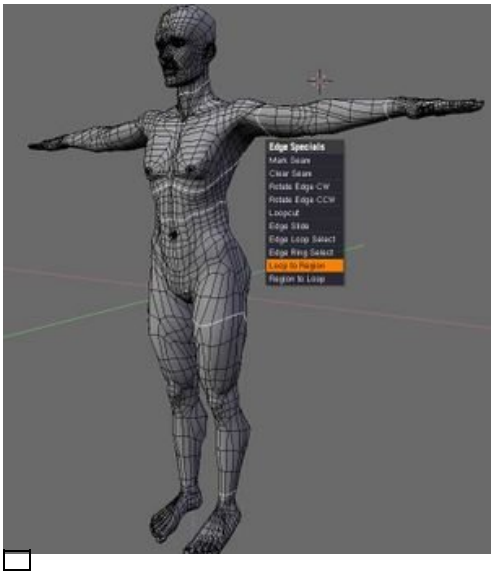
Examples: Loop to Region

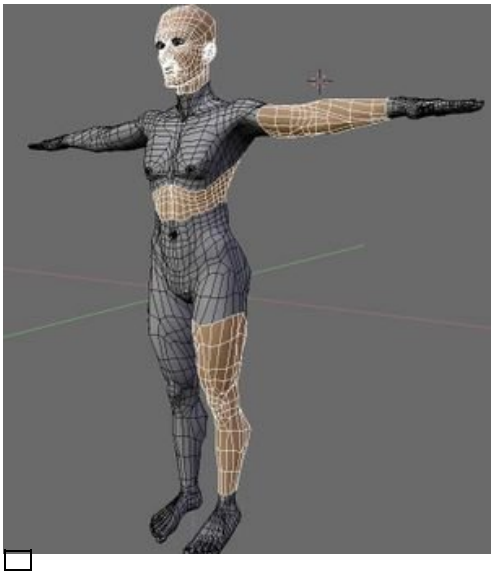


Selection.

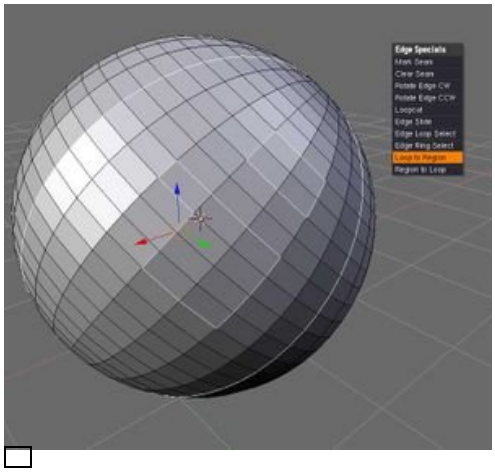


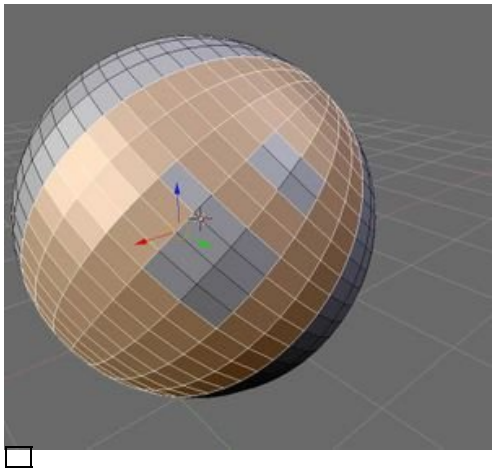
Loop to Region.





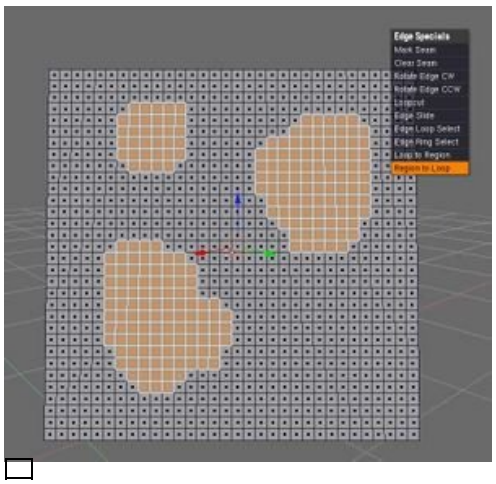
Selection. This tool handles multiple loops fine, as you can see.

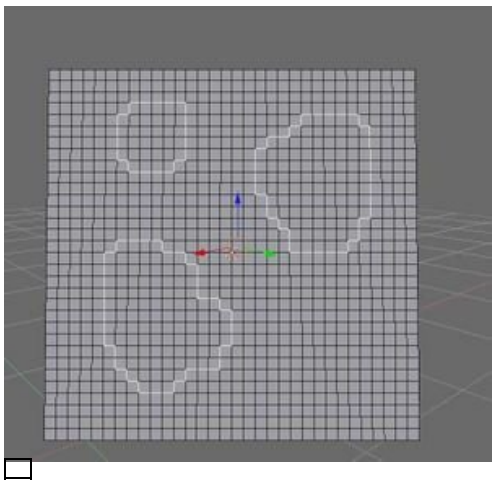




Selection. This tool handles â← holesâ← just fine as well.

Example: Region to Loop





Selection. This is the â← logical inverseâ← of loop to region.

Loop Subdivide

Mode: Edit mode (Mesh)

Hotkey: CtrlR

Menu: Mesh Â» Edges Â» Loop Subdivide...

Description

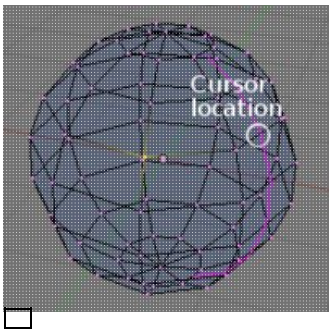
Loop Subdivide splits a loop of faces by inserting a new edge loop intersecting the chosen edge. The tool is interactive and has two steps:

1. Previsualising the cut

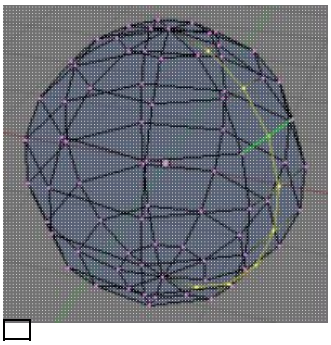
The cut to be made is marked with a magenta coloured line as you move the mouse over the various edges. In (*1. Pre-visualising the cut*), the mouse cursor was located where the white circle is located. This caused the loop line to appear at the mid point of the edge. The to-be-created edge loop stops at the poles where the existing face loop terminates.

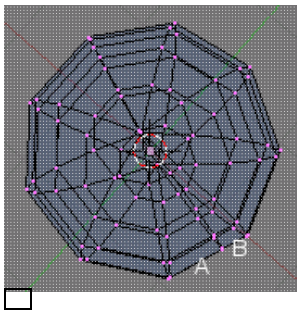
2. Sliding the new edge loop

Once an edge is chosen via LMB, that edge is highlighted in green (*2. Sliding the new edge loop*), and you can move the mouse along the edge to determine where the new edge loop will be placed. This is identical to the Edge Slide tool. Clicking LMB again confirms and makes the cut at the pre-visualised location, or clicking MMB forces the cut to exactly 50%.
(*3. Loop split edge completed*) shows the new faces and edges, â← Aâ← and â← Bâ← . The view is rotated so that the new faces and edges are clearly visible from the top of the sphere.



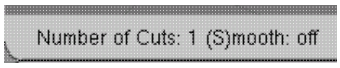
1. Pre-visualising the cut.


2. Sliding the new edge loop.



3. Loop split edge completed.

Options



Initial Loop Split header.

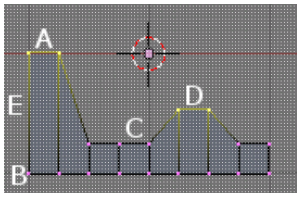
1. Previsualising the cut

- Upon initial activation of the tool 3D window header changes to show the Number of Cuts (*Initial Loop Split header*). Entering a number with the keyboard, scrolling Wheel or using + NumPad and - NumPad changes the number of cuts (maximum of 130). These cuts are uniformly distributed in the original face loop, and *you will not be able to control their positions* (no step â← 2â←).
- S changes the cut to â← smoothâ← mode. By default, new vertices for the new edge loop are placed exactly on the pre-existing edges. This keeps subdivided faces flat, but can distort geometry, particularly when using Subdivision Surfaces. If smooth mode is on then new vertices are not placed on the previous edge but shifted outwards by a given percentage, similar to the Subdivide Smooth command.

2. Sliding the new edge loop

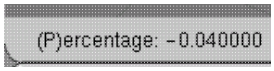
- P switches between *Proportional* and *Percentage* modes. The default mode is *Percentage*.
- In *Proportional* mode, Wheel, or â and â changes the selected edge for calculating a proportion.
- Holding Ctrl or â § Shift control the precision of the sliding. Ctrl snaps movement to 10% steps per move and â § Shift snaps movement to 1% steps. The default is 5% steps per move.

Examples



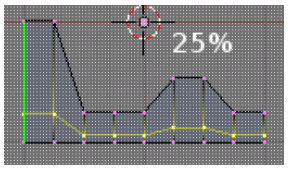
Loop Example Grid.
In order to explain *Proportional* and *Percentage* modes we can use a very simple mesh laid out like a 2Ã 9 grid, (*Loop Example Grid*). The vertices at â← Aâ← and â← Dâ← have been moved in order to emphasize the difference between the two modes. The vertices at the level â← Câ← and â← Bâ← remain unchanged. â← Eâ← is an area of interest when looking at *Proportional* mode.

Percentage mode

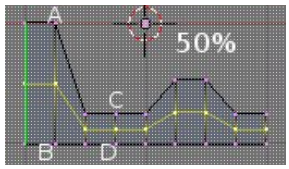


Percentage header.
In *Percentage* mode the 3D window header changes to (*Percentage header*) showing a number between **-1** and **1** with **0** representing 50% or midway between.

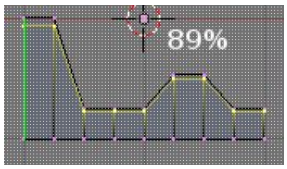
As you move the mouse the percentage changes and the edge loop line, drawn in yellow, moves as a percentage of the distance between the edge marked in green as shown in (*25% between*), (*Mid-way*) and (*89% between*).



25% between.



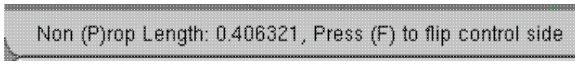
Mid-way.



89% between.

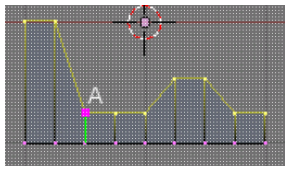
The yellow loop line is *always* the same percentage along edge of the edges that are being cut, regardless of the edgesâ← lengths. For example, in (*Mid-way*) the yellow loop line is exactly halfway between vertex â← Aâ← and â← Bâ← and it is also exactly halfway between vertex â← Câ← and â← Dâ← . For (*25% between*) you can see that the yellow line loop is *always* 25% along each of the cut edges.

Proportional mode

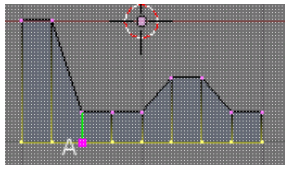


Proportional header.
Proportional face loop splitting keeps the shape of the newly cut edge loop the same as one of the edge loops that it cuts between, rather than cutting at a percentage along each perpendicular edge.

In *Proportional* mode the 3D window header changes to (*Proportional header*) showing the position along the length of the currently selected edge which is marked in green. Movement of the sliding edge loop is restricted to this length. As you move the mouse the length indicator in the header changes showing where along the length of the edge you are.



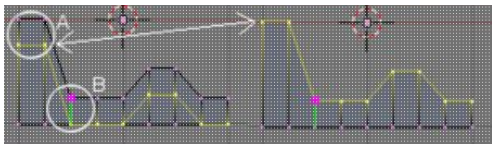
Vertex Marker.



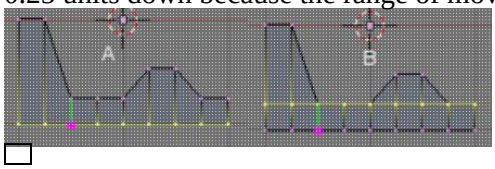
Vertex Marker Opposite.

Unlike *Percentage* mode, *Proportional* mode treats the edge as having a start and end vertex with the start marked by a magenta marker (*Vertex Marker*). The start vertex (â← Aâ←), can be flipped to the opposite vertex using F (*Vertex Marker Opposite*).

Moving the mouse moves the cut line towards or away from the start vertex, but the loop line will only move as far as the length of the currently selected edge, conforming to the shape of one of the bounding edge loops.



Proportional Range.
(*Proportional Range*) shows an example of how the distance is restricted by the length of the current edge (â← Bâ←). Looking at (â← Aâ←), you can see that the line loop has moved the same distance. If the line only moves 0.2 units on the selected edge then the line only moves 0.2 units everywhere else in the face loop region. The portion of the line loop at â← Aâ← hasnâ← t gone all the way to the â← bottomâ← because the selected edge is only 0.25 units in length. The line portion at â← Aâ← will not be able to move more than 0.25 units down because the range of movement is restricted to the length of the selected edge.



Proportional Range Flipped.
(*Proportional Range Flipped*) is another example where the start vertex has been flipped while using the same selected edge as compared to (*Proportional Range*). You can see that movement is still restricted to the length of the selected edge. The yellow edge loop line stays straight, conforming to the bottom bounding edge loop because the cut is placed a constant distance from the bottom edge loop, along the crossed edges.

Delete Edge Loop

Mode: Edit mode (Mesh)

Hotkey: X or Delete

Menu: Mesh Â» Edges Â» Delete Edge Loop

Description



Erase menu.
Delete Edge Loop allows you to delete a selected edge loop if it is between two other edge loops. This will create one face-loop where two previously existed.

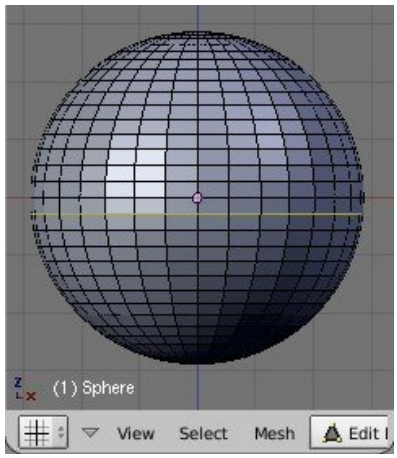
Note
The Edge Loop option is very different to the Edges option, even if you use it on edges that look like an edge loop. Deleting an edge loop merges the surrounding faces together to preserve the surface of the mesh. By deleting a chain of edges, the edges are removed, deleting the surrounding faces as well. This will leave holes in the mesh where the faces once were.

Limitations & Workarounds

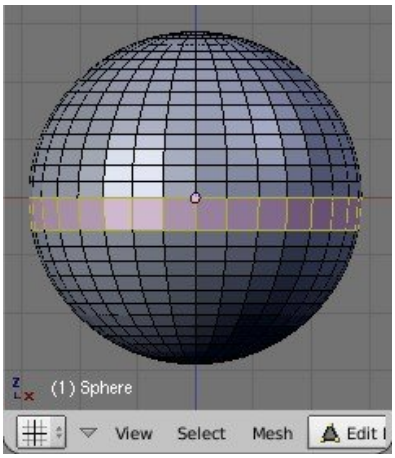
For Edge Loop Delete to work correctly, a single edge loop must be selected. The same restrictions as those of Edge Slide apply, see this page for more details.

Examples

The selected edge loop on the UV Sphere has been deleted and the faces have been merged with the surrounding edges. If the edges had been deleted by choosing Edges from the (*Erase Menu*) there would be an empty band of deleted faces all the way around the sphere instead.



Before Delete Edge Loop.



After Delete Edge Loop.

Knife Subdivide

Mode: Edit mode (Mesh)

Panel: Editing context (F9) â Mesh Tools

Hotkey: K or â § ShiftK

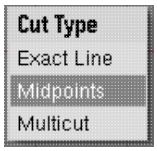
Menu: Mesh Â» Edges Â» Knife Subdivide...

Description

Knife Subdivide subdivides **selected** edges intersected by a user-drawn â← knifeâ← line. For example, if you wish to cut a hole in the front of a sphere, you select only the front edges, and then draw a line over the selected edges with the mouse. The tool is interactive, and only works with primary edges; selected either implicitly by selecting all or explicitly by box-selecting or â § ShiftRMB -clicking a few edges.

Tools donâ← t work on Modifiers
In general, you cannot use any tool on a mirrored side, as that side is just a mirror image of the primary side. Tools also do not work on subsurfed or multires â← edgesâ← shown; use the tool by working on vertices/edges/faces on the primary part of your mesh when using a modifier.

Options



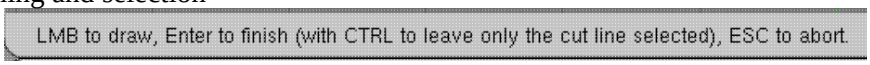
Knife Tool â Cut Type.
When you press K, the popup menu appears where you select the type of cut you can make:

- Exact Line divides the edges exactly where the knife line crosses them.
- Midpoints divides an intersected edge at its midpoint.
- Multicut makes multiple parallel cuts. An additional number input is presented, allowing you to select the number of cuts.

Drawing the cut line

When using Knife Subdivide, the cursor changes to an icon of a scalpel and the 3D View header shows (*Knife Tool 3DWindow header*). You can draw connected straight lines by clicking LMB and moving repeatedly or you can create *freehand* lines by pressing and holding LMB while dragging. Also, exact cuts on the vertices can be made by holding Ctrl while cutting. MMB constrains the cut line to a vertical or horizontal screen axis.

Confirming and selection



Knife Tool 3DWindow header.
Pressing Esc or RMB at any time cancels the tool, and pressing â µ Enter confirms the cut, with the following options:

- â µ Enter will leave selected every edge except the new edges created from the cut.
- Ctrlâ µ Enter will select only the new edges created from the cut. Note: only edges that intersect the hand-drawn selected edges will be selected.

Topology
: Knife subdivide uses the same options as the other subdivide tools, located in the Editing context. If the Beauty option is toggled selected faces are only subdivided along the longest 2 sides. If both Beauty and Short options are toggled, selected faces are only subdivided along the shortest 2 sides. Note: Using edge select mode to select only the edges you wish to subdivide creates a more accurate subdivision than using the Beauty toggle.

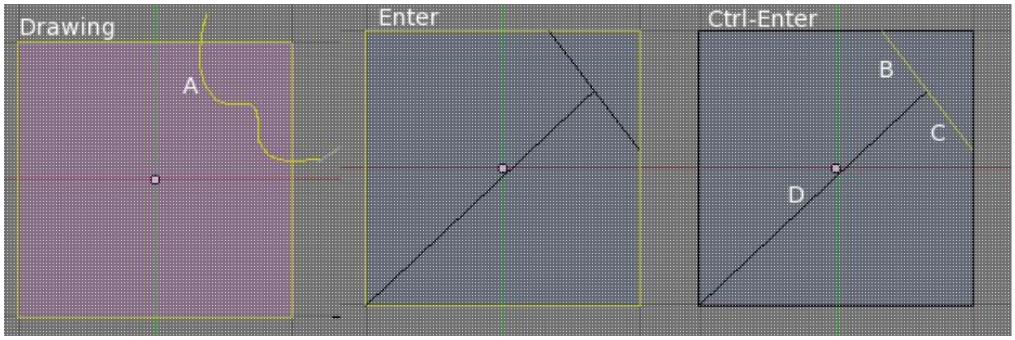
Examples

Exact Line Cut Type

(*Exact Line before and after*) is an example of using the Exact Line knife. The cut is determined from the hand-drawn line labeled â← Aâ← in the plane labeled â← Drawingâ← .

The plane labeled â← Enterâ← is the result of hitting the â µ Enter key. The intersections on the edges of the plane are where the drawn line actually intersects, no matter how wiggly the line is. In addition, all the edges have been selected other than the newly created edges from the cut tool itself.

The plane labeled â← Ctrl-Enterâ← is the result of hitting Ctrlâ µ Enter. In this case only the newly created edges, â← Bâ← and â← Câ← , are selected while edge â← Dâ← is not. â← Dâ← is a secondary edge added as a side effect of the cut tool.



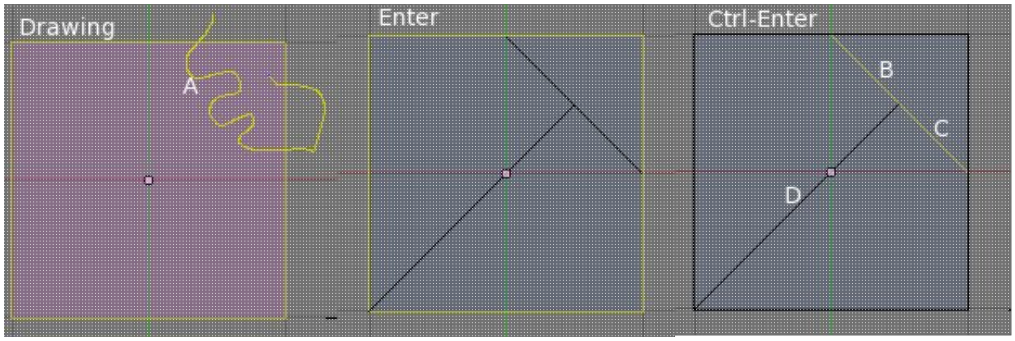
Exact Line, before and after.

Midpoints Cut Type

(*Midpoints before and after*) is an example of using the Midpoints knife. The cut is determined from the hand-drawn line labeled â← Aâ← on the plane labeled â← Drawingâ← . Notice how the line labeled â← Aâ← intersects the right edge twice; only the first intersection will be considered during the cut.

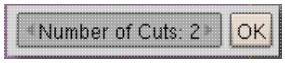
The plane labeled â← Enterâ← is the result of hitting â µ Enter. The intersections on the edges of the plane are the mid points of each edge, regardless of where the line was drawn. All the edges have been selected other than the newly created edges from the cut tool itself.

The plane labeled â← `Ctrl-Enter`â← is the result of hitting Ctrlâ µ Enter. In this case only the newly created edges, â← Bâ← and â← Câ← , are selected while edge â← Dâ← is not. â← Dâ← is a secondary edge as a result of the cut tool.



Midpoints before and after.

MultiCut Type

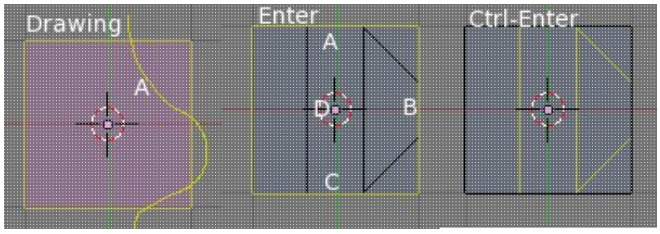


Number of Cuts.
This cut type presents a popup dialog asking for the Number of Cuts, which defines how many equally spaced cuts the tool should make for each intersecting edge. For example, the default of **2** generates two intersections or three new edges for each intersection of the hand-drawn line.

(*MultiCut before and after*) is an example of using the MultiCut knife. The cut is determined from the hand-drawn line (â← Aâ←) on the plane labeled â← `Drawing`â← , while using the default **2** as the number of cuts. The line was drawn so that it deliberately intersected three edges.

The grid labeled â← `Enter`â← is the result of hitting â µ Enter. There are two cuts equally spaced on each edge intersected by the hand-drawn line, labeled â← Aâ← , â← Bâ← , â← Câ← and â← Dâ← . The cut tool do not produce any secondary edge here.

The grid labeled â← `Ctrl-Enter`â← is the result of hitting Ctrlâ µ Enter. In this case only the newly created edges are selected.



MultiCut before and after.

Limitations & Workarounds

The cut lines can be drawn with any number of segments, but only one intersection is detected per edge. Crossing back over an edge multiple times does not perform additional cuts on it. Snap to grid is not currently implemented, but is being looked at for future releases.

Optimizations

With a large mesh, it will be quicker to select a smaller number of vertices, those defining only the edges you plan to split since the Knife will save time in testing selected vertices for knife trail crossings.

Snap to Mesh

Mode: Edit or Object modes

Hotkey: â § Shiftâ Tab

Menu: Mesh / Object Â» Transform Â» Snap

Description

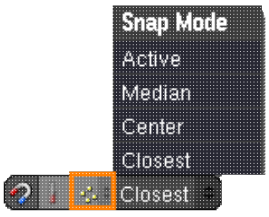
Used in conjunction with Grab, Extrusion, Scale or Rotation modes, this tool allows you to snap objects (in Object mode) or vertices/edges/faces (in Edit mode) to a target vertex/edge/face.

Once the tool is activated you are ready to drag your element(s) to its(their) destination. Move your mouse pointer to the target where you want to snap and hold down Ctrl, move your pointer to adjust (especially if your target is an edge or face).

When satisfied with the result, confirm with LMB (as usual, you can cancel everything with a RMB clic or by hitting Esc).

Options

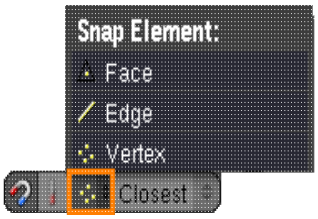
Snap Mode



Snap Mode menu.

- Closest: move the closest point of the selection to the target.
- Center: move the current transformation center to the target. Can be used with 3D cursor to snap with an offset.
- Median: move the median of the selection to the target.
- Active: move the active element (vertex in Edit mode, object in Object mode) to the target.

Snap Element



Snap Element menu.

- Vertex: the target will be a vertex.
- Edge: the target will be an edge.
- Face: the target will be a face.

The following shows different Element and Mode options: Download Demo Video (theora).

Align Rotation



Snap align.
Align Objectâ← s Z axis with the normal of the target element.

Normals are interpolated between the two vertex normals when snapping to edges, otherwise, face normals and vertex normals are directly used.

Only works with Translations (Grab) done in Object mode.

The following video shows a tree being snapped and aligned on a ground mesh: [Download Demo Video] (theora).

Vertex Groups

A mesh is a set of connected Vertices, sometimes thousands of vertices for the more complex objects. Blender allows you to group these vertices for several main reasons:

- Re-using parts of a mesh for making copies.
- Hiding â← everything elseâ← while you work on details.
- Documentation and explanation to others.
- Armatures deformation.
- Generating particles from only the group.
- Controlling the velocity of particles emitted.

Armatures
Vertex Groups can be automatically created for each bone in an armature. However, that process is pretty involved and for more information on Armatures and Bone Vertex Groups, click here. The rest of this section will focus on user-defined vertex groups.

Why use Vertex Groups?

Vertex groups identify sub-components of an object, like the legs of a chair or the hinges of a door. By â← bookmarkingâ← such regions in vertex groups you can easily select and work on them in isolation without having to create separate objects. With the hide function you can even remove everything else from view.

Vertex groups also make it easy to cull out and duplicate a part of the mesh many times. Consider modeling a Lego block. The most simple block consists of a base and a nipple. To create a four-nipple block, you would want to be able to easily select the nipple vertices, and, still in Edit mode, duplicate them and position them where you want them.

Another use for vertex groups is for skinning an armature. If you want to animate your mesh and make it move, you will define an armature which consists of a bunch of invisible bones. As a bone moves, it deforms or moves the vertices associated with it. Not all of the vertices, but some of them; the ones assigned to it. So, when you move the bone â← Armâ← , the arm bone moves the â← Armâ← vertices, and not the â← Legâ← vertices. In this way, parts of the mesh can stretch and move around, while other parts remain stationary.

By selecting the group in the Vertex group set of controls of the Extras panel (Particle sub-context, Object context), the weight painting of the group might control how much particles come out, of which size, at which velocity, etc. Recall that hair is a type of particle system, so define a vertex group called â← `Scalp`â← and use it to tell Blender to emit hair from the scalp!

Another great use for vertex groups is for keeping track of vertex selections. For instance, I was modeling a coin-like object with a raised, beveled border. For some reason, I couldnâ← t do a loop select around the base of the border, so I had to carefully select the vertices. I knew Iâ← d need to use that same selection several times in the next few minutes while working on the model, so I named and saved the group, saving me a lot of work later.

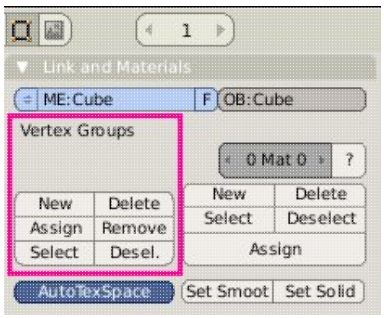
Vertex Groups Basics

Creating a Vertex Group

Mode: Edit mode

Panel: Editing context (F9) â Link and Materials Â» Vertex Groups Â» New

Hotkey: CtrlG Â» Add Selected to New Group



Vertex Group default panel.
By default, an object does not have any groups, and all of its vertices are hanging out there in cyberspace as loners. The image to the right highlights the Vertex Groups buttons in hot pink. These buttons are located in a Buttons window in the Editing context (F9) on the Link and Materials panel. They are shown when an object with vertices is selected **and** being edited (â Tab). You can tell when an object is in Edit mode because your 3D window cursor is a cross-hair.

Groups are only for Vertices
Vertex Groups are only available for objects that have vertices. Text objects, for example, cannot have vertex groups and the panel is not shown when that kind of object is selected. Vertex Groups are only shown when an object with vertices (a mesh) is being edited.

To create a vertex group, LMB click the New button. When you do, a new vertex group (named, surprisingly, â← `Group`â←) is created, and the panel shows you a Weight numeric slider/entry/scroll box. Any selected vertices are not yet assigned to the new vertex group, you must click the Assign button to actually allocate vertices to the newly created vertex group. Note that using the Shortcut CtrlG Â» Add Selected to New Group, you do all this in one step.

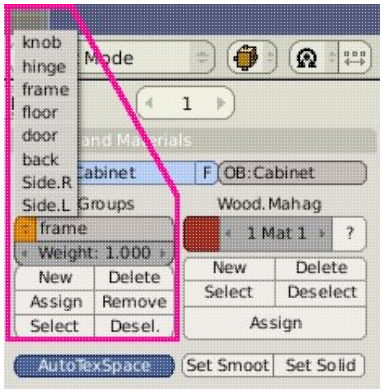
Check Your Assignment
Itâ← s a good idea to make sure the vertices have been properly assigned to the group by using the deselect and select buttons. If nothing happens, just hit the Assign button to add the selected vertices to the group.

Naming Vertex Groups

Mode: Edit mode

Panel: Editing context (F9) â Link and Materials Â» Vertex Groups â dropdown list

To name a group something other than the creative â← Groupâ← , â § ShiftLMB click the name field, and type in the name you want.



Cabinet Vertex Group example.
For example, consider the model of a kitchen cabinet. The cabinet consists of three vertical walls (two sides and a back), a floor and countertop, a door frame, a door, a knob and two hinges. You may or may not, at some point, to be able to model the door opening. You may want to make the cabinet a single door or later easily modify it to be a double door (with two knobs). You may wish to copy the knob design and use it for the drawers which you will be modeling later. In this case, you would want to define at least three vertex groups: â← `Base`â← , â← `Door`â← , and â← `Knob`â← . If you were writing a user manual, you would want your example to contain each possible group for maximum re-use and selection, as shown.

Access the group list by clicking the list selector button next to the group name. Select a group by clicking on any named group.

Assigning Vertices to a Group

Mode: Edit mode

Panel: Editing context (F9) â Link and Materials Â» Vertex Groups Â» Assign

Hotkey: CtrlG Â» Add Selected to Active Group

To add vertices to a group you do the following:

1. Select the group you want to work with from the group list.
2. Use your mouse to â § ShiftRMB select more vertices that you want in that group.
3. LMB click the Assign button, or make CtrlG Â» Add Selected to Active Group.

Keep in mind that a vertex can be assigned to multiple groups.

Note
When using the Assign button to assign selected vertices to a vertex group, any vertices that were already in the vertex group are not removed, so Assign adds extra vertices to the selected vertex group.

Seeing a Vertex Group

Mode: Edit mode

Panel: Editing context (F9) â Link and Materials Â» Vertex Groups Â» Select

From experience, we have found that is is best to start first by seeing the existing vertices in a group, before adding more or removing some. To do this, de-select all vertices by pressing A once or twice in the 3D window, until the User Preferences header shows â← `Ve:0-x`â← , where â← *x*â← is the number of vertices in your mesh. This means that zero (`0`) vertices are selected. The Vertex count is located just to the right of the Blender Version.

Then, with the appropriate group active, press the Select button. In your 3D window, the vertices that belong to the active group will be selected and highlighted.

Removing Vertices from a Group

Mode: Edit mode

Panel: Editing context (F9) â Link and Materials Â» Vertex Groups Â» Remove

Hotkey: CtrlG Â» Remove Selected from Active Group / Remove Selected from All Groups

To remove vertices from a group:

1. Select the group you want to work with from the group list.
2. Select the vertices you want to remove from the vertex group.
3. LMB click the Remove button, or use CtrlG Â» Remove Selected from Active Group.

Note you can remove selected vertices from all groups at once, using the Vertex Groups popup menu (CtrlG Â» Remove Selected from All Groups).

Deselecting Vertices

Mode: Edit mode

Panel: Editing context (F9) â Link and Materials Â» Vertex Groups Â» Deselect

Sometimes you will want to see if any vertices are still loners. To do so, select All the vertices in the 3D window. For each Vertex Group, LMB click the Deselect button to de-select the vertices in that group. Repeat the de-selection for each group. When you run out of groups, any vertices left highlighted are the loners. Sort of like picking baseball teams.

Deleting a Group

Mode: Edit mode

Panel: Editing context (F9) â Link and Materials Â» Vertex Groups Â» Delete

To delete a vertex group, select it from the list and click Delete. Yes, itâ← s as simple as that. Any vertices that belonged to that group are unassigned from that group. However, please keep in mind that vertices can belong to many Groups. When they are unassigned from one group, they still belong to their other groups.

Using Vertex Groups in Practice

Assume you have defined the groups used in our cabinet example. Here are some examples of common things you might want to do involving Vertex Groups.

Duplicating Parts

You now want to make that cabinet a double door model:

1. Select the cabinet object (RMB) and enter Edit mode (â Tab).
2. Ensure that **no** vertices are selected (`Ve:0-x` â← remember?).
3. Select the â← `Knob`â← vertex group from the dropdown menu.
4. Click the Select button.
5. Move your mouse into the 3D window.
6. Duplicate that sub-mesh by pressing â § ShiftD. The vertices are copied, selected, and grabbed.
7. Move the mouse over to position the new knob.
8. LMB to drop the sub-mesh.

The duplicated vertices belong to the same group(s) as their originals.

To assign this new knob to its own group, click New, name it something like â← `Knob.L`â← and click on Assign. See Creating a Vertex Group and Assigning Vertices to a Group.

Left and Right naming convention
Certain features of Blender can perform related actions on groups that are left and right counterparts of each other. If you end a name in â← `.L`â← or â← `.left`â← and its counterpart â← `.R`â← or â← `.right`â← , Blender may be able to easily mirror its actions for you.

You can read more about the naming convention in Editing Armatures: Naming conventions. The convention for armatures/bone applies here as well.

Simplifying a Vertex Group

You may have correctly surmised that the original â← `Knob`â← group now has both sets of vertices: the original and the duplicated ones. Youâ← ve created a â← `Knob.L`â← group, but there is no corresponding â← rightâ← group. The â← `Knob`â← group really needs to be renamed and contain only the vertices for the right knob. To correct this,

1. Ensure the new â← `Knob.L`â← vertices are still selected (the ones that donâ← t belong).
2. Select the original â← `Knob`â← Vertex Group from the list.
3. Click the Remove button.

To test your work, deselect all vertices and click the Select button. Only the vertices from the original knob should highlight. Rename this group â← `Knob.R`â← .

Repeat the above for the â← `Door`â← and â← `Hinge`â← group, and you now have a two-door cabinet model. Note that you will have to either make the doors narrower or cabinet wider to accommodate the new door.

Combining Groups

To create a â← Knobsâ← (plural group), you could:

1. Ensure that no vertices are selected.
2. Select the â← Knob.Lâ← group (select its name from the list and click Select).
3. Select the â← Knob.Râ← group (ditto).
4. Observe that selecting one set of vertices does not deselect the others; the selection process adds on vertices to the selection.
5. Click the New button, and name the group â← Knobsâ← .

Focus on a part of your model

You want to make an inset panel on the door. To work on the door sub-mesh without cluttering up your screen with all the other vertices, you would:

1. Ensure that **all** vertices are selected (you can use A for that).
2. Deselect the â← Doorâ← group by selecting its name from the Vertex Group list and clicking Deselect, leaving everything **but** the door selected.
3. With your cursor in the 3D window, Hide the selected vertices. Poof! They disappear.
4. AltH makes them visible again.

Separating a part into its own

Now, the patent lawyer calls and says that you must patent your hinge design to keep anyone else from copying it; you need to separate the hinge out from the cabinet mesh:

1. Ensure that **no** vertices are selected.
2. Select the â← Hingeâ← vertices (select the name from the Vertex Group list, and click Select).
3. With your cursor in the 3D window, seParate them into their own object.
4. The remaining cabinet vertices are left. â Tab out of Edit mode and RMB click the floating hinge object. Note that it is conveniently called â← Cabinet.001â← , and has all the same Vertex Groups as the original. Delete those groups you do not need, rename the object â← Hingeâ← .

Parent it to the original (and now hinge-less) â← Cabinetâ← object (include the parent by â § ShiftRMB clicking the cabinet, and pressing CtrlP). Now, when you move your cabinet, the hinges move with it.

Finding Groups for a Vertex

As you are rigging and animating the deformation of a mesh, you might need to find out which groups a vertex belongs to, and to adjust the weights of influence each group has on that particular vertex.

1. Select the vertex.
2. In Edit mode, press N to open the Transform Properties window for that vertex.
3. Open the drop-down menu that shows all vertex groups to which it belongs.
4. From this window you also can assign weights to the vertex for each group.

About Weight

By default, every vertex in a group has a weight of 1.00. If a vertex belongs to multiple groups, it has a combined weight. When influenced by a bone or other object, it is moved by an amount proportional to its weight; heavier vertices move less. So, a middle vertex belonging to two groups (each with a weight of 1.00) would move half as much as a left vertex that only belonged to one group. This weighting system provides realistic deformation of a mesh when bones move, for example, around the shoulder area, where some of the vertices belong to both the chest and the arm groups.

You can set the weight of all vertices in a group using the Weight numeric control. For more advanced weighting, please read Weight Painting. Weight Painting allows you to smoothly blend individual vertex weights so that meshes deform smoothly.

Weight paint

The Weight Paint Mode is used to create and modify Vertex Groups. A vertex may not only be a member of one or more Vertex Groups, but also may have a certain weight in each group. The weight symbolises its influence on the result.

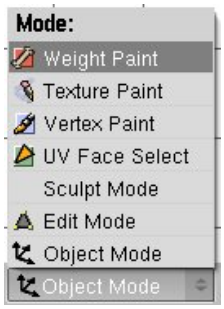


Image 1: Changing to Weight Paint Mode.
Mode: Object Mode

Hotkey: Ctrlâ Tab

Menu: Mode Menu (***Image 1***)

When you change to Weight Paint Mode you see the selected object (if you have not created any groups yet), in a slightly shaded, blue colour (***Image 2***).

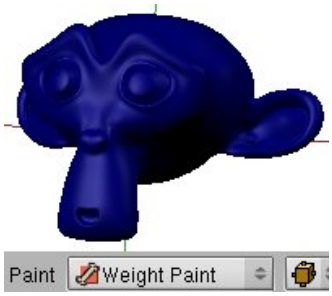


Image 2: An object in Weight Paint-Mode.
The colour visualises the weight of each vertex of the currently active group. A vertex drawn in blue indicates either: a weight of zero, not in the active group, or not in any group at all.

You assign the weight of each vertex by painting on the mesh with a certain colour. Starting to paint on a Mesh will automatically create a new Vertex Weight Group (when no group existed or is active). If a vertex doesnâ← t belong to the active group it is automatically added (if the option Vgroup is not set), even if you paint with a weight of â← 0â← . The used colour spectrum is shown in (***Image 3***).

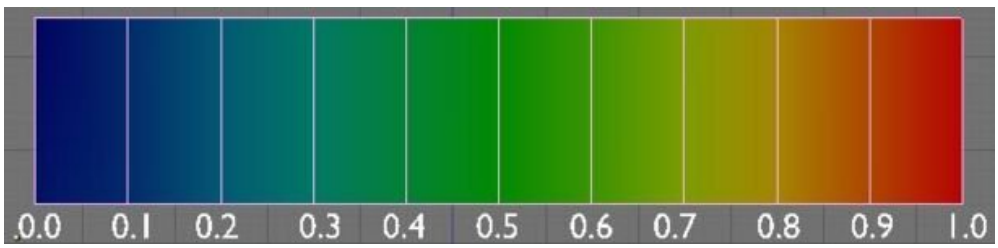


Image 3: The colour spectrum and their respective weights.

You paint on the mesh with a brush. The colour only influences the vertices, neither the faces nor the edges. So donâ← t try and paint these. There is a tool panel for the brush in the Editing-Buttons (F9) as well as in the 3D Window (press N to open it).

Weight-Painting survival tips:

A few tips that will save you some hassle when painting weight:

- Press F in Weight Paint Mode to activate the Face Select mode. Now you can select faces to paint and Hide faces like in Edit Mode.
- Press B to Border-select faces to paint using LMB. Use the RMB border select to exclude the selected faces from painting.
- Draw a *Clipping Border* with AltB. It will separate a visible part of the 3D-window. You can draw only in this part. If you press AltB again the *Clipping Border* will be removed.
- Turn off Soft in the Paint-Panel. If you have Soft activated you will reach the target value only after several repeated paint actions, and itâ← s especially difficult to reach **0.0**.

Paint-Panel

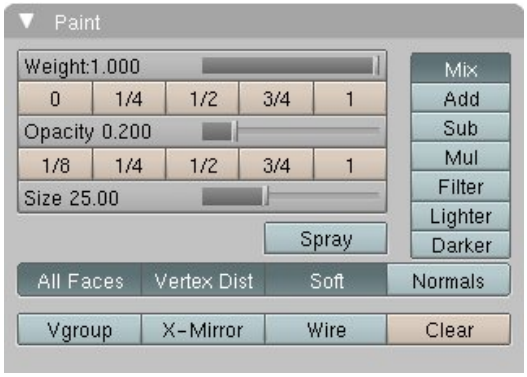


Image 4: The Paint-Panel in the Editing-Buttons.
The tools in the Paint panel are sophisticated, and you can apply weight in the finest nuances. But normally you wonâ← t need all these options, and you will apply weight using a few techniques. The most used and important settings are **drawn in bold**.

- **Weight:** The weight (colour) that is used to paint. The button row below contains five weight presets to paint. By default, painting works with an absolute fixed amount (like Gimp or Photoshop defaults), so you can set for example â← 0.2 weightâ← to vertices while keeping the mouse button pressed, whatever their original weight was.
- **Opacity:** The extent to which the vertex colour changes while you are painting.
- Size: The size of the brush, which is drawn as a circle during painting.
- Spray: The Spray option accumulates weights on every mouse move.
- Mix/Add/â←¦: The manner in which the new colour replaces the old when painting.
 - Mix: The new colour replaces the old colour. Applying more and more of the new colour will eventually turn it the new colour.
 - Add: The new colour is added to the old. Note that you must think in weight here (not in RGB colours): adding blue (**0.0**) to something wonâ← t modify it, adding green (**0.5**) to green will give red (**1.0**), â←¦
 - Sub: The new colour is subtracted from existing. Here again you have to think in terms of â← weightâ← , and not â← RGB coloursâ← .
 - Mul: The new colour is multiplied by the old.
 - Filter: Paint based on alpha value.

 - ♦ Lighter: Only paints vertices â← darkerâ← (lower weight) than the current â← colourâ← , â← lighteningâ← them.
 - ♦ Darker: Only paints vertices â← lighterâ← (higher weight) than the current â← colourâ← , â← darkeningâ← them.
- All faces: If this is turned off, you will only paint vertices which belong to a face on which the cursor is. This is useful if you have a complicated mesh and you would paint on visually near faces that are actually quite distant in the mesh.
- Vertex Dist: Paints only on vertices which are actually under the brush. If you switch this off, all vertices belonging to faces touched by the brush are painted. If you have a sparse mesh and use subsurfaces you want to keep this on.
- **Soft:** This specifies that the extent to which the vertices lie within the brush also determine the brushâ← s effect. Itâ← s extremely difficult to paint with zero then. You want to turn this off in all normal situations.
- Normals: The vertex normal (helps) determine the extent of painting. This causes an effect as if painting with light.
- **Vgroup:** Only vertices which belong to the active vertex group are painted. Very useful for clearing up and refining vertex groups without messing other groups up.
- **X-mirror:** Use the X-Mirror option for mirrored painting on groups that have symmetrical names, like with extension .R/.L, or _R/_L. If a group has no mirrored counterpart, it will paint symmetrical on the active group itself. You can read more about the naming convention in Editing Armatures: Naming conventions. The convention for armatures/bones apply here as well.
- Wire: Show additionally the wireframe of the objects. Since it shows the subsurfaced wire itâ← s quite useless. Itâ← s better to use the Select-Mode (see below).
- **Clear:** Removes all vertices from the active group.

Tools

If you have a complex mesh it is nearly impossible to reach all vertices in Weight Paint mode. And it is quite difficult to tell where the vertices are exactly. But thereâ← s a very good and easy solution: the Select Mode.

Mode: Weight Paint Mode

Hotkey: F

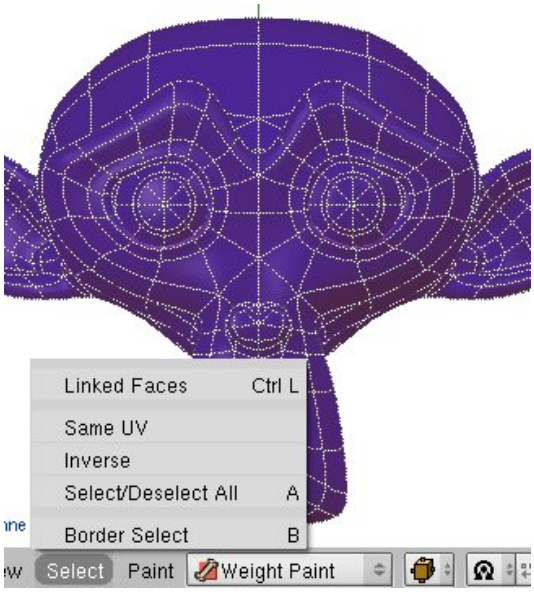


Image 5: Select-Menu in Weight Paint-Mode.
Select Mode has many advantages over the normal Weight Paint mode:

1. The original mesh is drawn, even when subsurface is active. You can see the vertices you have to paint over.
2. You can select faces, only the vertices of selected faces are painted on.
3. Selecting tools include:
 - ♦ RMB â← Single faces. Use â § ShiftRMB to select multiple.
 - ♦ A â← All faces, also to deselect.
 - ♦ B â← Block/Box selection.
 - ♦ B-B â← Select with brush.
 - ♦ CtrlL â← Select linked.
 - ♦ In the Select-Menu: Faces with Same UV, also invert selection (Inverse).
4. You may hide selected faces with H and show them again with AltH (***Image 6***).

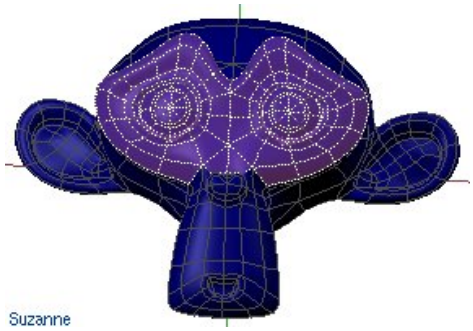


Image 6a: Select interfering facesâ←¦

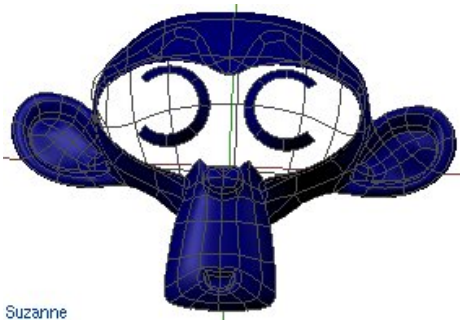


Image 6b: â←¦ and hide them with H.

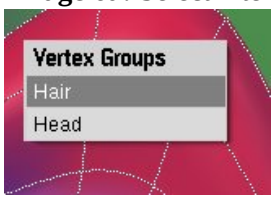


Image 7: A vertex belonging to two vertex groups.

- To constraint the paint area further you may use the *Clipping Border*. Press AltB and LMB -Drag a rectangular Area. The rest of the 3D window is hidden. To show everything again, just hit another time AltB.
- If you want to know which groups a Vertex belongs use â § ShiftLMB . You can change between the groups in the appearing Pop-Up Menu (***Image 7***).
- N in the 3D-window opens a Paint-panel instead of the transform properties panel (***Image 8***).

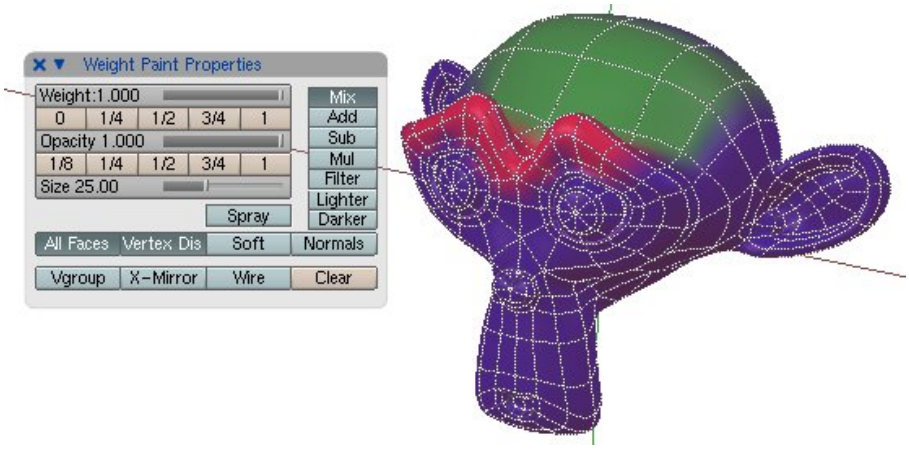


Image 8: Weight Paint Properties Panel in the 3D-Window.

Scripts in *Paint*-Menu

Weight Gradient
: This script is used to fill the selected faces with a gradient (***Image 3*** & ***Image 9***). To use the script, select the faces you wish to apply the gradient to. Click twice on the mesh to set the start and end points of the gradient. The colour under the mouse will be used for the start and end blend colours, so you have to weight paint first. Holding â § Shift or clicking outside the mesh on the second click will blend the first colour to nothing.

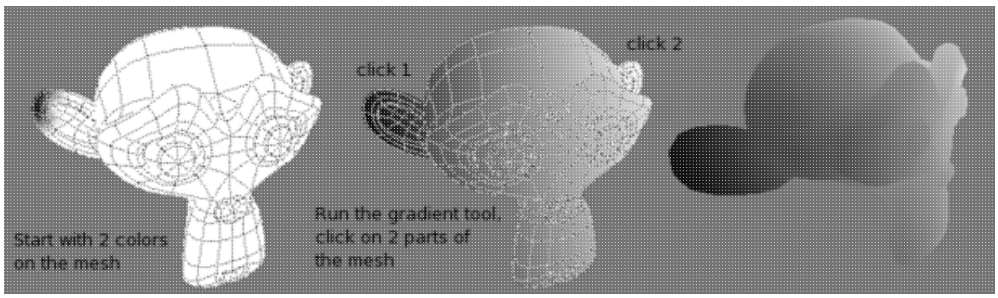


Image 9: Setting the gradient and result.

Normalise/Scale Weight
: Maximizes weights to a user set peak and optionally scales other groups too to keep the proportion of the weights even.

Grow/Shrink Weight
: Uses the mesh topology to expand/contract the vertex weights (works like colour bleeding).

Clean Weight
: Removes vertices with low weights from the current group.

Weight-Painting for Bones

This is probably the most often used application of weight painting. When a bone moves, vertices around the joint should move as well, but just a little, to mimic the stretching of the skin around the joint. Use a â← lightâ← weight (10-40%) paint on the vertices around the joint so that they move a little when the bone rotates. While there are ways to automatically assign weights to an armature (see the Armature section), you can do this manually. To do this from scratch, refer to the process below. To modify automatically assigned weights, jump into the middle of the process where noted:

- Create an Armature.
- Create a Mesh that will be deformed when the armatureâ← s bone(s) move.
- With the Mesh selected, create an Armature modifier for your mesh (located in the Editing context, Modifiers panel). Enter the name of the Armature.

(Pick up here for modifying automatically assigned weights)

- Select the Armature in 3D View, and bring the armature to **Pose mode** (Ctrlâ Tab, or the 3D View window header mode selector).

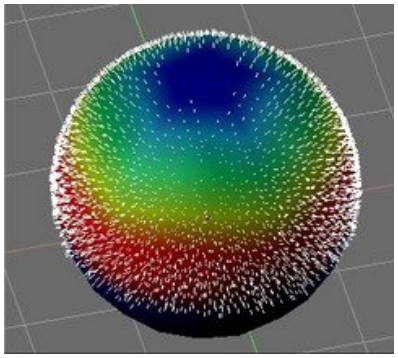
- Select a desired bone in the Armature.
- Select your mesh (using RMB) and change immediately to Weight Paint mode. The mesh will be coloured according to the weight (degree) that the selected bone movement affects the mesh. Initially, it will be all blue (no effect).
- Weight paint to your heartâ← s content. The mesh around the bone itself should be red (generally) and fade out through the rainbow to blue for vertices farther away from the bone.

You may select a different bone with RMB . If the mesh skins the bones, you will not be able to see the bones because the mesh is painted. If so, turn on X-Ray view (Buttons window, Editing context, Armature panel). While there on that panel, you can also change how the bones are displayed (Octahedron, Stick, B-Bone, or Envelope) and enable Draw Names to ensure the name of the selected bone matches up to the vertex group.

If you paint on the mesh, a vertex group is created for the bone. If you paint on vertices outside the group, the painted vertices are automatically added to the vertex group.

If you have a symmetrical mesh and a symmetrical armature you can use the option X-Mirror. Then the mirrored groups with the mirrored weights are automatically created.

Weight-Painting for Particles



Faces or vertices with zero weight generate no particles. A weight of 0.1 will result in 10% of the amounts of particles. This option â← conservesâ← the total indicated number of particles, adjusting the distributions so that the proper weights are achieved while using the actual number of particles called for. Use this to make portions of your mesh hairier (is that really a word?) than others by weight-painting a vertex group, and then calling out the name of the vertex group in the VGroup: field (Particles panel, Object context).

Subdivision Surfaces

Mode: Any mode

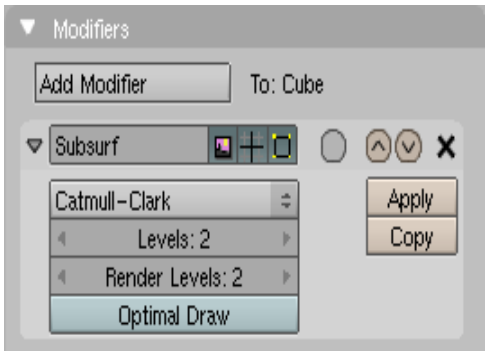
Panel: Modifiers (Editing context, F9)

Hotkey: â § ShiftO (toggle subsurf in Object mode)

Description

Subdivision Surface is a method of subdividing the faces of a mesh to give a smooth appearance, to enable modeling of complex smooth surfaces with simple, low-vertex meshes. This allows high resolution mesh modeling without the need to save and maintain huge amounts of data and gives a smooth *organic* look to the object. With any regular mesh as a starting point, Blender can calculate a smooth subdivision on the fly, while modeling or while rendering, using a simple dummy subdivision surface (Subsurf in short), or the smarter Catmull-Clark one.

Options



Modifiers panel.
Subsurf is a modifier. To add it to a mesh, press Add Modifier and select Subsurf from the list.

Subsurf Type
This drop-down list allows you to choose the subdivision algorithm:
- ◊ Simple Subdiv. â← Just subdivides the surfaces, without any smoothing (similar to Levels time W â Subdivide, in Edit mode). Rarely useful!
- ◊ Catmull-Clark â← Default option, subdivides and smooths the surfaces.

Levels
Defines the display resolution, or level of subdivision *for Blender 3D views*.

Render Levels
This is the subdivision level *used during rendering*.

Note
These two settings allow you to keep a fast and lightweight approximation of your model when interacting with it in 3D, but use a higher quality version when rendering.

Optimal Draw

Restricts the wireframe display to only show the original mesh cage edges, rather than the subdivided result, to help visualization.

Subsurf UV
: When enabled, the UV maps will also be subsurfed (i.e. Blender will add â← virtualâ← coordinates for all sub-faces created by this modifier).

To view and edit the results of the subdivision (â← isolinesâ←) while youâ← re editing the mesh, you must enable the Editing Cage mode by clicking in the gray circle in the modifier panel header (next to the arrows for moving the modifier up and down the stack). This lets you grab the points as they lie in their new subdivided locations, rather than on the original mesh.

Hints

You can use â § ShiftO if you are in Object mode to switch Subsurf On or Off in *both* the Object mode *and* at rendering. To only turn the subsurf view on/off in Object mode (to reduce lag), press â § ShiftAltO. The Subsurf level can also be controlled via Ctrl1 to Ctrl4, but this only affects the visualization sub-division level.

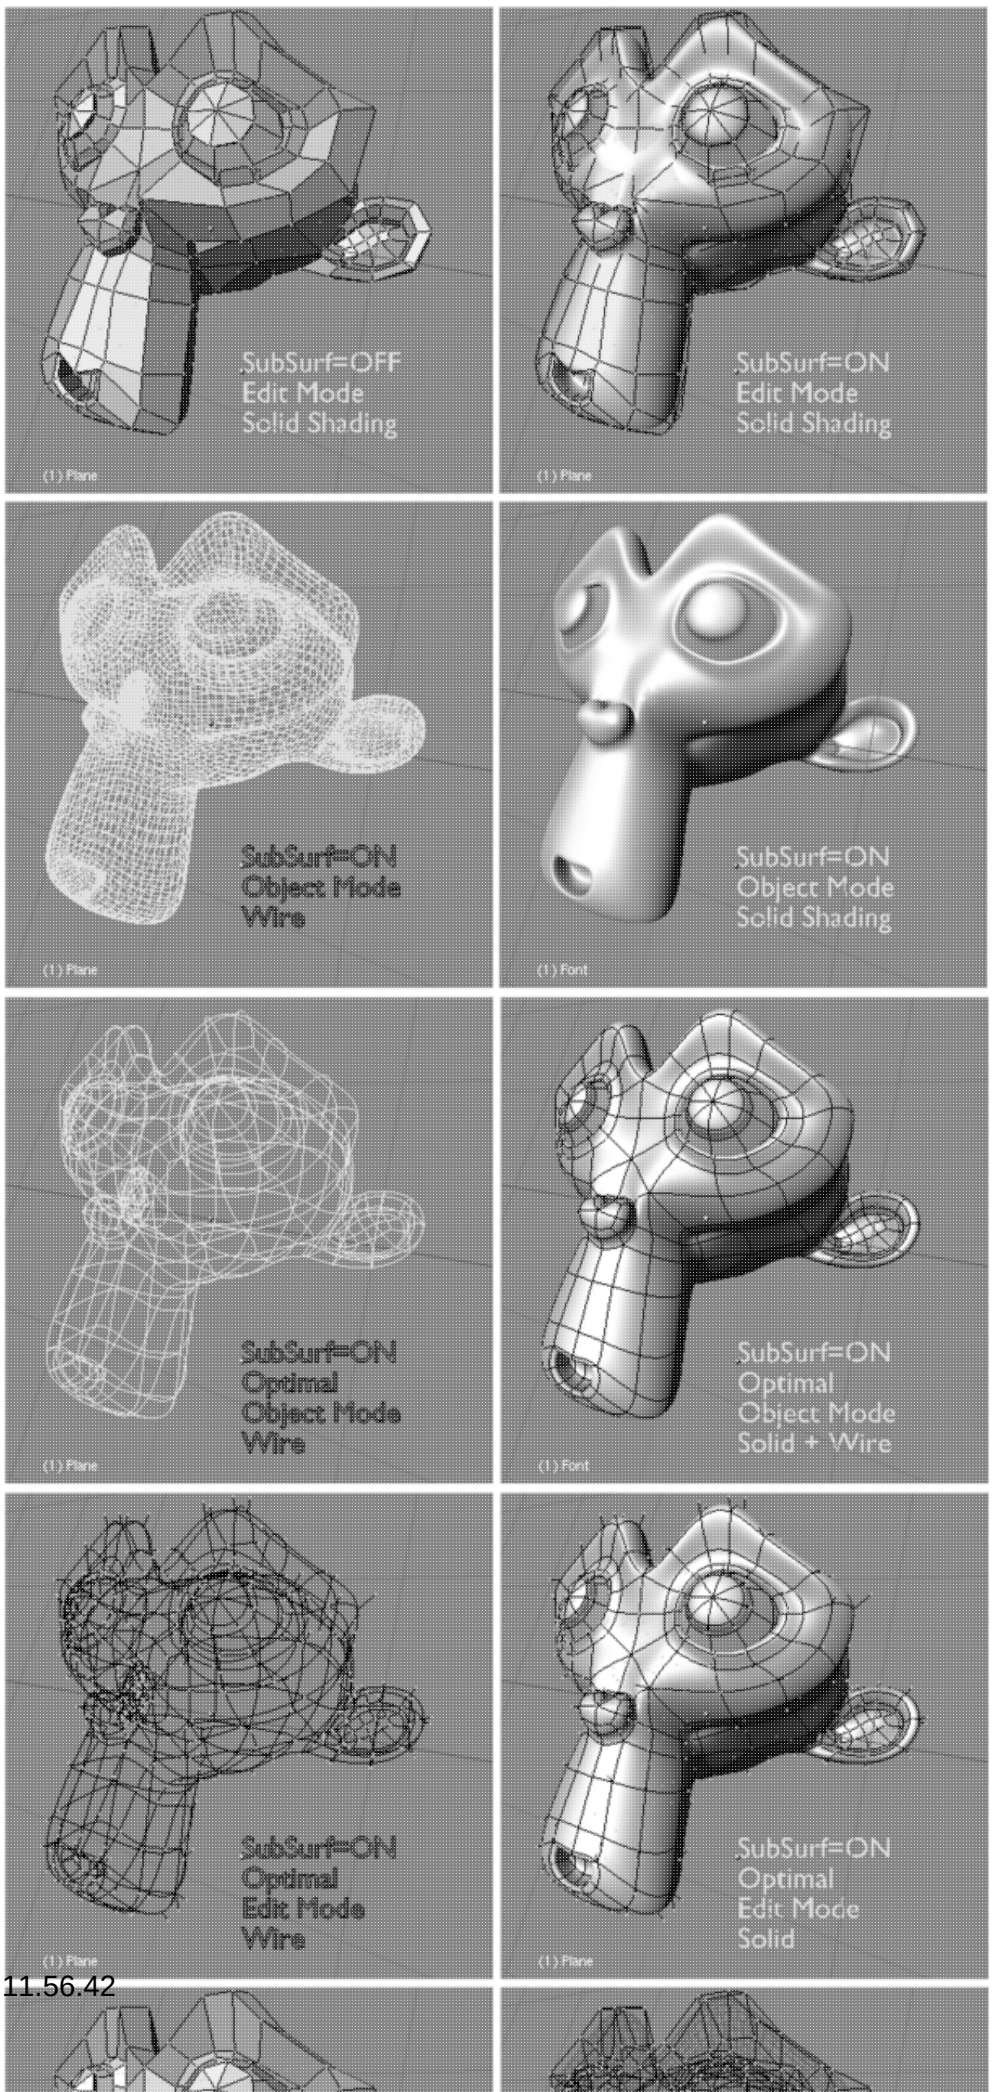
A subsurfed mesh and a NURBS surface have many points in common such as they both rely on a â← coarseâ← low-poly â← meshâ← to define a smooth â← high definitionâ← surface, however, there are notable differences:

- NURBS allow for finer control on the surface, since you can set â← weightsâ← independently on each control point of the control mesh. On a subsurfed mesh you cannot act on weights (except for edges crease, see bellow).
- Subsurfs have a more flexible modeling approach. Since a subsurf is a mathematical operation occurring on a mesh, you can use all the modeling tools and techniques described in the chapter on mesh modeling. There are more techniques, which are far more flexible, than those available for NURBS control polygons.

Since Subsurf computations are performed both real-time, while you model, and at render time, and they are CPU intensive, it is usually good practice to keep the subsurf Levels low (but non-zero) while modeling, and higher while rendering (Render Levels).

Examples

(*Subsurfed Suzanne*) is a series of pictures showing various different combinations of Subsurf options on a Suzanne mesh.


SubSurf=OFF
Edit Mode
Solid Shading
(1) Plane
SubSurf=ON
Edit Mode
Solid Shading
(1) Plane
SubSurf=ON
Object Mode
Wire
(1) Plane
SubSurf=ON
Object Mode
Solid Shading
(1) Font
SubSurf=ON
Optimal
Object Mode
Wire
(1) Plane
SubSurf=ON
Optimal
Object Mode
Solid + Wire
(1) Font
SubSurf=ON
Optimal
Edit Mode
Wire
(1) Plane
SubSurf=ON
Optimal
Edit Mode
Solid
(1) Plane

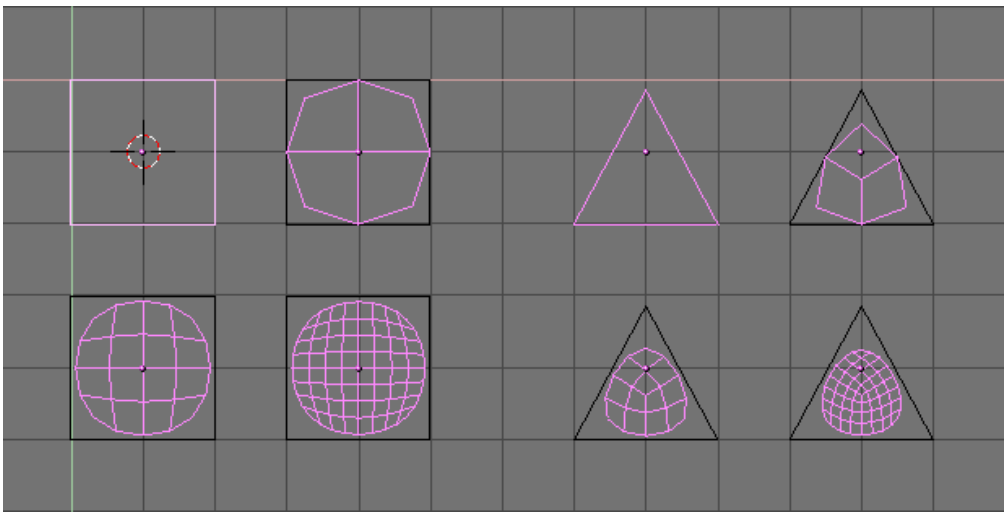
Subsurfed Suzanne.

(*SubSurf of simple square and triangular faces*) shows a **0**, **1**, **2**, **3** level of subsurf on a single square face or on a single triangular face. Such a subdivision is performed, on a generic mesh, for *each* square or triangular face.

It is obvious that the number of faces produced by the n level of subdivision is:

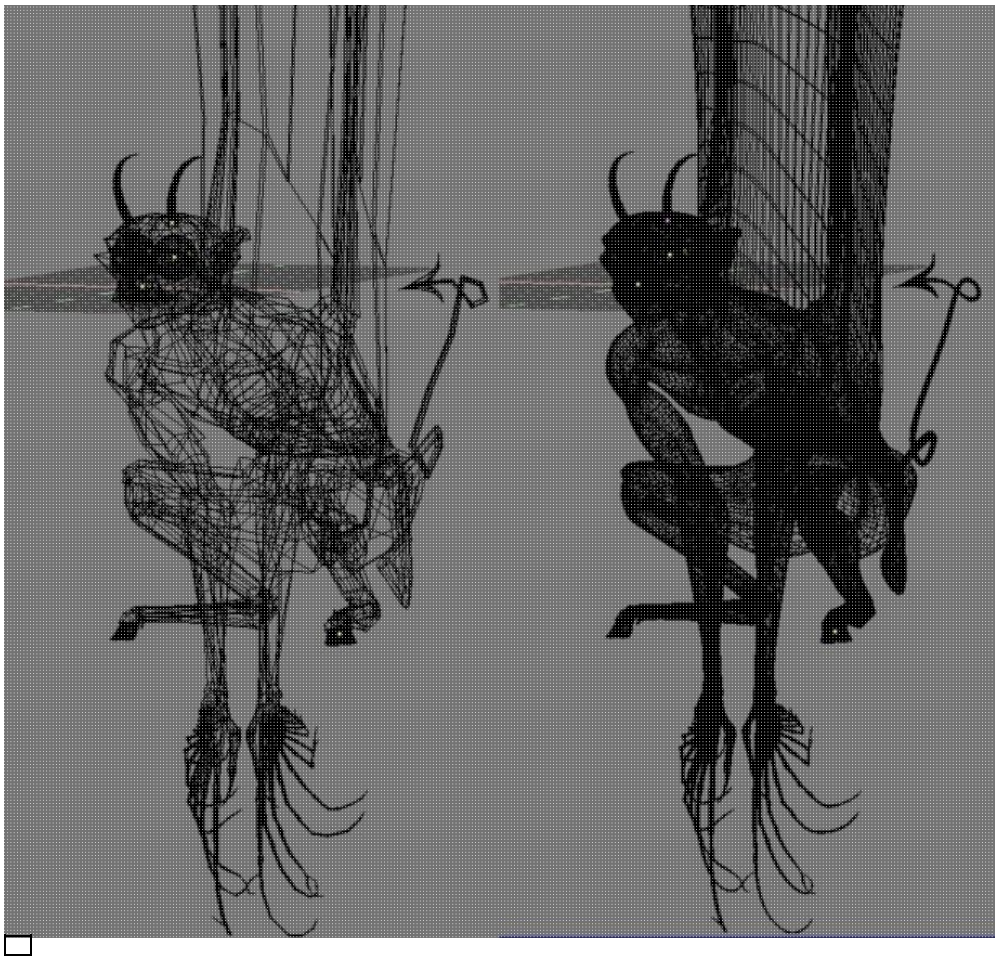
- For a quadrangle : 4^n.
- For a triangle : 3 Ã 4^{n-1}.

This dramatic increase of face (and vertex) number results in a slow-down of all editing, and rendering, actions and calls for lower Subsurf level in the editing process than in the rendering one.

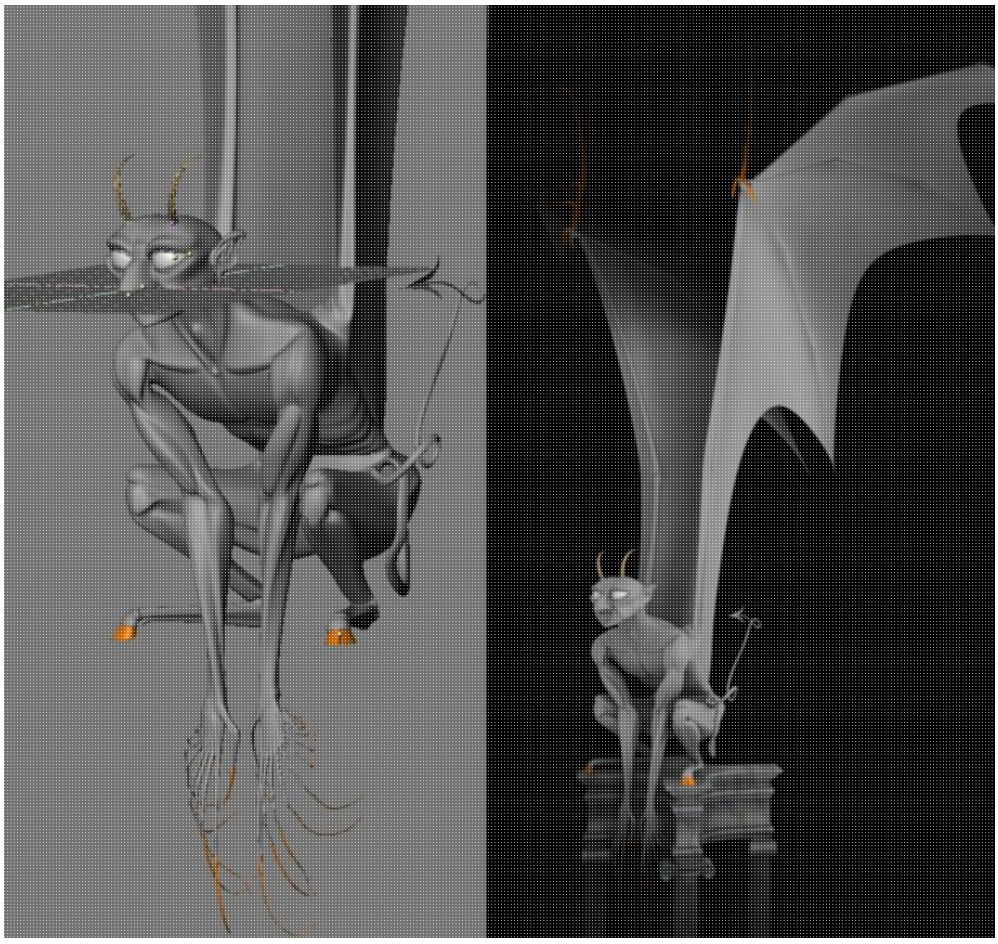


SubSurf of simple square and triangular faces.

The Subsurf tool allows you to create very good â← organicâ← models, but remember that a regular mesh with square faces, rather than triangular ones, gives the best results. (*A Gargoyle base mesh (left) and pertinent level **2** subsurfed mesh (right)*) and (*Solid view (left) and final rendering (right) of the Gargoyle*) show an example of what can be done with Blender subsurfs.



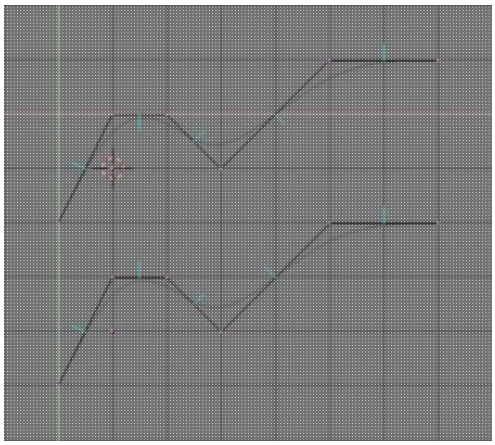
A Gargoyle base mesh (left) and pertinent level **2** subsurfed mesh (right).



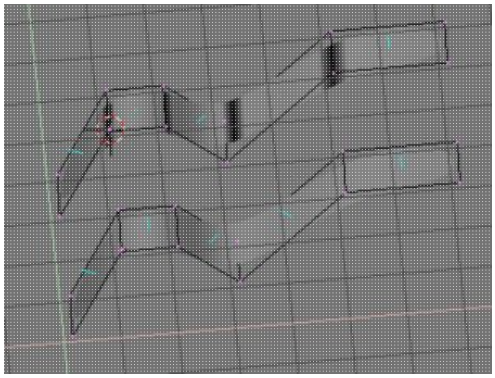
Solid view (left) and final rendering (right) of the Gargoyle.

Limitations & Workarounds

Blenderâ← s subdivision system is based on the Catmull-Clark algorithm. This produces nice smooth subsurfed meshes but any subsurfed face, that is, any small face created by the algorithm from a single face of the original mesh, shares the normal orientation of that original face.



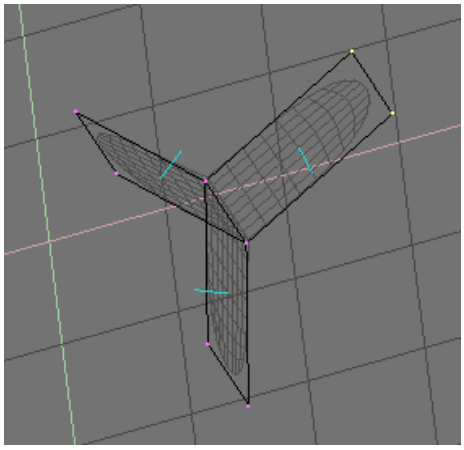
Side view of subsurfed meshes. With random normals (top) and with coherent normals (bottom).
This is not an issue for the shape itself, as (*Side view of subsurfed meshes. With random normals (top) and with coherent normals (bottom)*) shows, but it is an issue in the rendering phase and in solid mode, where abrupt normal changes can produce ugly black lines (*Solid view of subsurfed meshes with inconsistent normals (top) and consistent normals (bottom)*).



Solid view of subsurfed meshes with inconsistent normals (top) and consistent normals (bottom).
Use the CtrlN command in Edit mode, with all vertices selected, to recalculate the normals to point outside.

In these images the face normals are drawn cyan. You can enable drawing normals in the Editing context (F9), Mesh Tools More panel.

Note that Blender cannot recalculate normals correctly if the mesh is not â← Manifoldâ← . A â← Non-Manifoldâ← mesh is a mesh for which an â← outâ← cannot unequivocally be computed. From the Blender point of view, it is a mesh where there are edges belonging to *more* than two faces.



A â← Non-Manifoldâ← mesh.
(*A â← Non-Manifoldâ← mesh*) shows a very simple example of a â← Non-Manifoldâ← mesh. In general a non-manifold mesh occurs when you have internal faces and the like.

A non-manifold mesh is not a problem for conventional meshes, but can give rise to ugly artifacts when subsurfed. Also, it does not allow decimation, so it is better to avoid them as much as possible.

Use these two hints to tell whether a mesh is â← Non Manifoldâ← :

- The recalculation of normals leaves black lines somewhere.
- The Decimator tool in the Mesh panel refuses to work stating that the mesh is â← Non-Manifoldâ← .

Weighted creases for subdivision surfaces

Mode: Edit mode (mesh)

Panel: 3D View â Transform Properties (N)

Hotkey: â § ShiftE

Menu: Mesh Â» Edges Â» Crease Subsurf

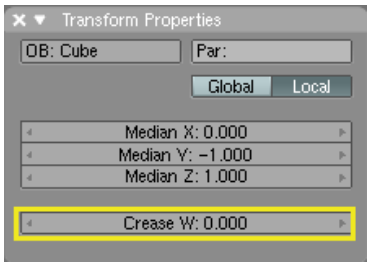
Description

Weighted edge creases for subdivision surfaces allow you to change the way Subsurf subdivides the geometry to give the edges a smooth or sharp appearance.

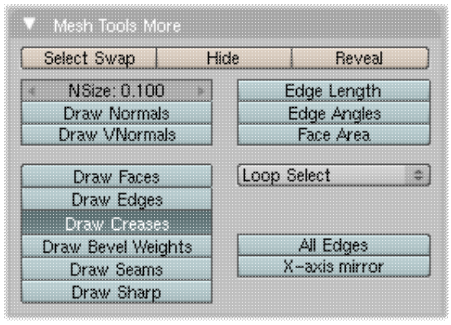
Options

The crease weight of selected edges can be changed interactively by using â § ShiftE and moving the mouse towards or away from the selection. Moving the mouse away from the edge increases the weight. You can also use Transform Properties (N) and enter the value directly. A higher value makes the edge â← strongerâ← and more resistant to subsurf. Another way to remember it is that the weight refers to the edgeâ← s sharpness. Edges with a higher weight will be deformed less by subsurf. Recall that the subsurfed shape is a product of all intersecting edges, so to make the edges of an area sharper, you have to increase the weight of all the surrounding edges.

You can enable an indication of your edge sharpness by enabling Draw Creases. See (*Mesh Tools More panel*).



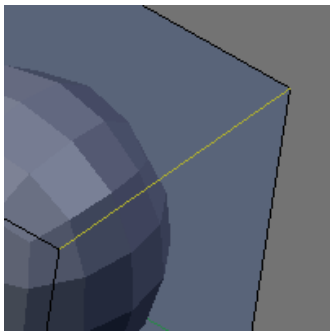
Transform Properties panel.

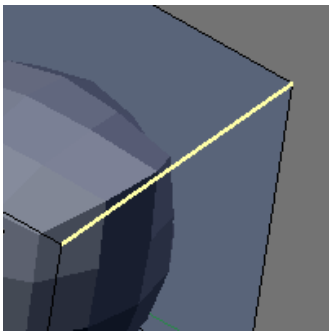


Mesh Tools More panel.

Examples

The sharpness value on the edge is indicated as a variation of the thickness on the edge. If the edge has a sharpness value of **1.0**, it will be thicker than if it had a sharpness value of **0.0**.


Crease: **0.0**.


Crease: **1.0**.

Multiresolution Mesh

Mode: Edit Mode, Object Mode

Panel: Editing Context (F9) â Multires

Description



Multires Monster head sculpt by Tom Musgrove.
Multires stands for 'multiple resolution mesh'. Multires allows you to edit the mesh at both high and low levels of complexity. Changes you make at one level of resolution propagate to all other levels.

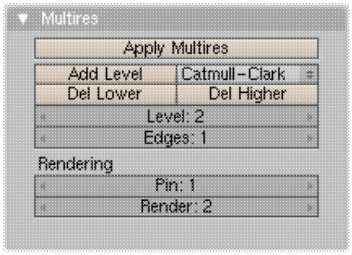
This feature is most often used for Sculpt Mode, where multires levels are added to a base mesh to sculpt with increasingly fine detail. While sculpting, viewing a high-complexity mesh is CPU-intensive. Multires allows the user to view the mesh at a lower level for positioning, lighting, and animating. Even when viewing a mesh at a low level, all the fine detail still remains in the multires data of the mesh and can be viewed at any time by setting the Level: value to a higher number.

Options

In the Multires panel press Add Multires. Then press Add Level to increase the levels of multires. Level 1 is the same as a mesh without multires. If you press Apply Multires it will convert the mesh to whatever the current level is that is selected. You can add multiple levels until you reach your RAM limit.



Multires panel, no multires data added yet.



Multires panel.

- Add multires - Initializes the mesh to accept multires levels.
- Apply multires - Appears only after adding multires to a mesh. Removes all the multires data, *leaving the mesh at the currently set resolution*.
- Add Level - Changes to the highest level and subdivides the mesh. This also modifies the vertices in all lower levels to match those in the highest level.
- Del Lower - Applies the changes to the mesh at the current level and removes all lower mesh levels.
- Del Higher - Deletes all multires levels above the current level.
- Edges - Determines the maximum level of edges (the highlighted outline of a mesh) drawn. If you use a solid shaded drawing, Wire must be enabled in the Draw panel (Object context, F7). Setting Edges to a high level will make the edges of the mesh closely follow the complexity of the mesh. This is especially useful for sculpting when you need to see the effect of your brush along the edge of a mesh.
- Pin - If you have a modifier on the mesh, this determines what level the modifier is applied during rendering. Any multires level above the modifier is disabled.
- Render - This determines what level of multires the model is rendered at. By default set to the highest available level.

Limitations

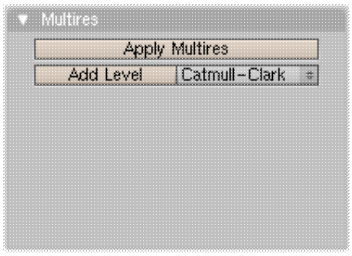
- Only the shape and not the topology of the mesh can be changed with multires enabled. Thus any tool that changes the topology (deleting or adding of faces) is deactivated, and trying to use one will raise the error message â← `Unable to complete action with multires enabled.`â← . This includes tools like extrusion, duplications, subdivisionsâ←¦ You can add modifiers, but you wonâ← t be able to apply them if they modify the topology.
- Multires currently incompatible with shapekeys.
- Some modifiers can result in a slow down of display and interaction if multires is also enabled.
- The low resolution mesh may change shape when switching to a higher resolution and back. This is because multires tries to fit the low level shape as good as possible to the high level shape, which is quite useful when sculpting a complex shape from a simpler one, but may not be desirable in all situations.

Workflow



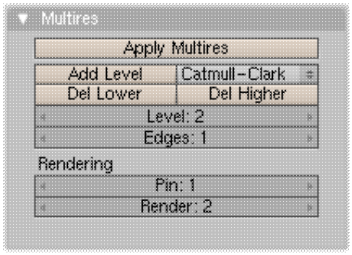
Multires panel, no multires data added yet.

- Select a mesh. In the Multires panel of the Editing context, click the Add Multires button. This simply sets up the mesh for multiple resolutions. It does not add any levels.

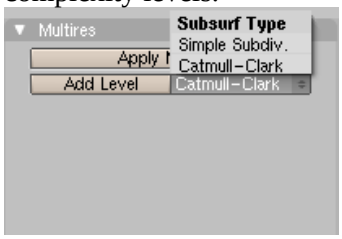


Multires data added. Click Add Level to add the first level of resolution.

- Click Add Level to add the first level of multiresolution data.



Multires panel after clicking the Add Level button once. Click "Add Level" at this point to add higher complexity levels.



Simple subdivision or Catmull Clark? Decisionsâ←¦

- With the first level of multires added, more options become available (see above). Level 1 is the same as the mesh without any extra resolution. Level 2 is the real first subdivided level. To add more levels with higher resolution, click Add Level again. You can choose between simple subdivision or Catmull Clark with the drop-down list. Simple subdivisions are useful when you want your mesh to keep it's existing shape without being smoothed.
- Adding more levels will increase the load on your CPU and make frame rate drop considerably. When you don't need to view the mesh in high detail, set the Level: value to Level 1 or 2.

- Enable Sculpt Mode and sculpt at the level that gives you a balance of performance and control.

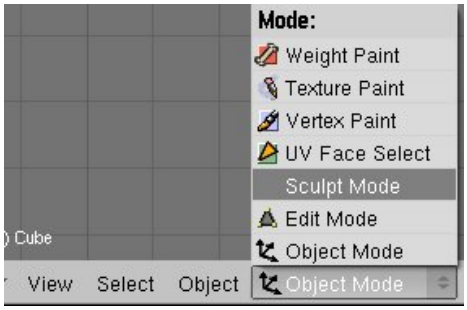
See also

- Dev:Source/Blender/Modeling/Multires - Development details and some examples.

Overview

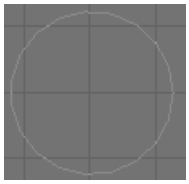
Sculpt Mode is similar to Edit Mode in that it is used to alter the shape of a model, but Sculpt Mode uses a very different workflow: instead of dealing with individual elements (vertices, edges, and faces), an area of the model is altered using a brush. In other words, instead of selecting a group of vertices, Sculpt Mode automatically selects vertices based on where the brush is, and modifies them accordingly.

Sculpt Mode



Sculpt Mode Dropdown.
Sculpt mode is selected from the mode menu of the 3D View header.

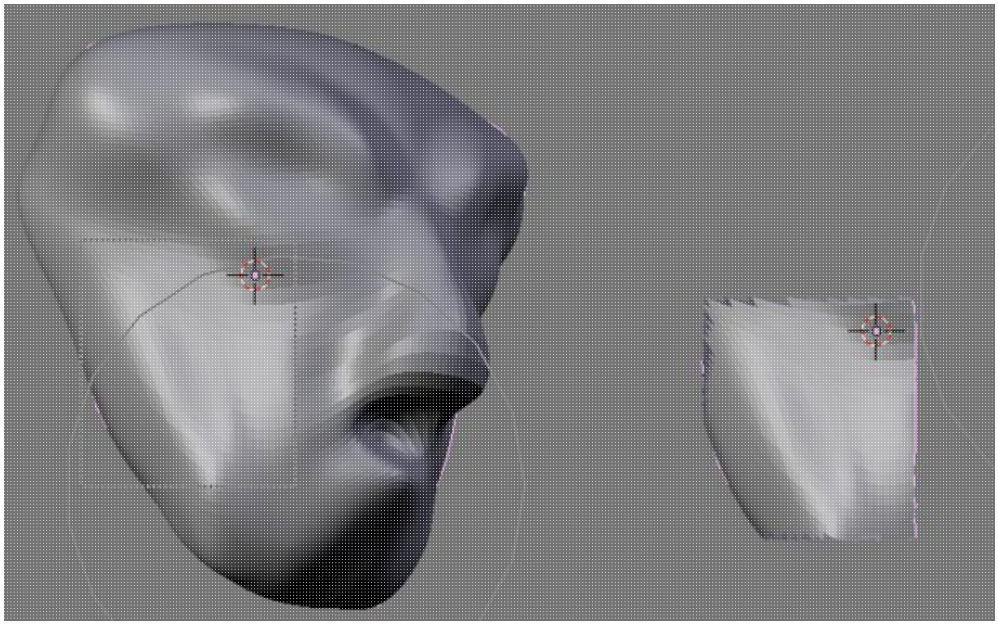
Once sculpt mode is activated a sculpt menu will appear in the 3d view header, and three tabs for Sculpt, Brush and Texture panels will appear in the Multires panel, in Editing context (F9). Also the cursor will change to a circle with a cross hair in the centre.



The cursor in Sculpt Mode.

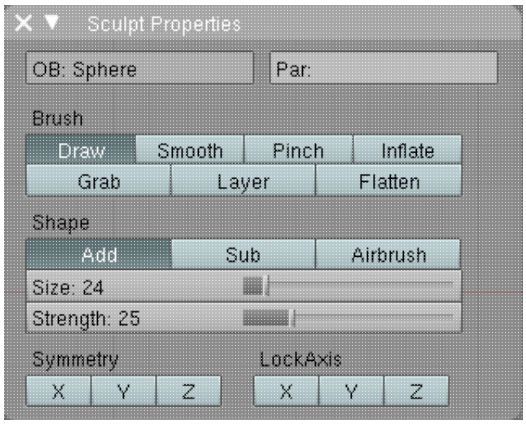
Hiding and Revealing Mesh

To hide a part of a mesh, Ctrlâ § ShiftLMB -drag around the part you want to keep - everything else will be hidden. Or Ctrlâ § ShiftRMB -drag to hide only the selected rectangle. To reveal all hidden parts, just hit AltH or click and release Ctrlâ § ShiftLMB .



Before and after Hiding.

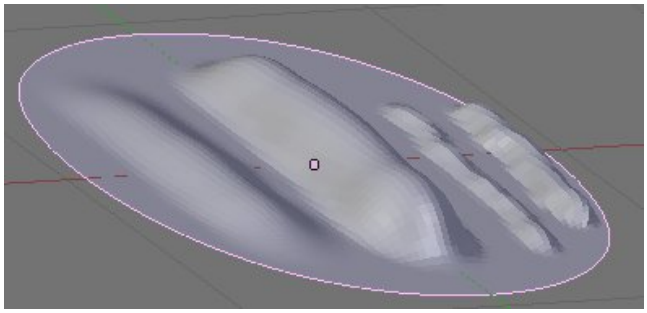
Sculpt Properties Panel



The Sculpt Properties panel.
This panel is opened with N or through the 3D view menu: Sculpt â Sculpt Properties. It is the exact equivalent of the Sculpt tab in the Editing context (F9).

Brushes

Sculpt Mode has seven brushes that each operate on the model in a unique way:



Drawing in various sizes and strengths.

Draw (D)
: Draws a smooth curve on the model following the brush; vertices are displaced in the direction of the average normal of the vertices contained within the brush.

Smooth (S)
: As the name suggests, eliminates irregularities in the area of the mesh within the brushâ← s influence.

Pinch (P)
: Pinch pulls vertices towards the centre of the brush. If Sub is active instead of Add, vertices are pushed away from the centre of the brush.

Inflate (I)
: Similar to Draw, except that vertices in Inflate mode are displaced in the direction of their own normals.

Grab (G)
: Grab is used to drag a group of points around. Unlike the other brushes, Grab does not modify different points as the brush is dragged across the model. Instead, Grab selects a group of vertices on mousedown, and pulls them to follow the mouse. The effect is similar to moving a group of vertices in Edit mode with proportional-editing enabled, except that Grab can make use of other Sculpt Mode options (like textures and symmetry.)

Layer (L)
: This brush is similar to Draw, except that the height of the displacement layer is capped. This creates the appearance of a solid layer being drawn. This brush does not draw on top of itself; brush stroke intersects itself. Releasing the mouse button and starting a new stroke will reset the depth and paint on top of the previous stroke.

Flatten (T)
: The Flatten will lower the height of the part of the mesh being worked on. Simply put, the direction of the flattening depends on the way the surface normals in the mesh are facing.

Modifiers

Brush Shape

Par
: The parent objectâ← s name.

Add and Sub
: Add causes the brush to pull an area of the model in the positive direction, Sub in the negative direction. This applies to following brushes: Draw, Inflate and Layer â← you have already seen the Pinch brush case (Add pulls vertices inward and Sub pushes vertices outward). Smooth, Grab and Flatten do not use these settings. Interactive toggling of brush direction is with holding down â § Shift. Or V can be used to toggle it until it is toggled again.

Airbrush (A)
: When enabled, this option causes the brush to continue modifying the model after mouse down without moving the mouse. If disabled, the brush only modifies the model when the brush changes its location. This option is not available for the Grab brush.

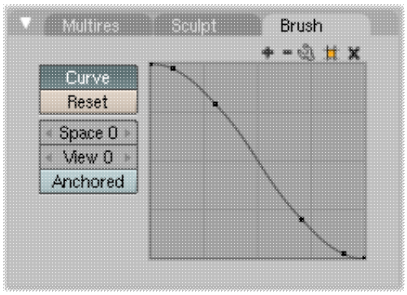
Size
: This option controls the radius of the brush, measured in pixels. F in the 3D view allows you to change the brush size interactively by dragging the mouse and then left clicking (The texture of the brush should be visible inside the circle). Typing a number then enter while in F sizing allows you to enter the size numerically.

Strength
: Strength controls how much each application of the brush affects the model. For example, higher values cause the Draw brush to add depth to the model more quickly, and cause the Smooth brush to smooth the model more quickly. This setting is not available for Grab.If the range of strengths doesnâ← t seem to fit the model (for example, if even the lowest strength setting still makes too large of a change on the model) then you can scale the model (in Edit Mode, not Object Mode). Larger sizes will make the brushâ← s effect smaller, and vice versa. You can change the brush strength interactively by pressing â § ShiftF in the 3D view and then moving the brush and then left clicking. You can enter the size numerically also while in â § ShiftF sizing.

Symmetry
: Mirror the brush across the selected local axes. Note that if you want to alter the directions the axes point in, you must rotate the model in Edit Mode, not Object Mode. Can be toggled via X, Y, and, Z respectively.

LockAxis
: These three buttons allow you to block any modification/deformation of your model along selected local axes, while you are sculpting it.

Brush and Texture Panels



The Brush panel.

Sculpt Mode can take full advantage of the range of options offered by Blenderâ← s texture system. Brush textures are accessed using a similar interface to that used by the Material or World context: there are nine

texture slots in the Sculpt Mode Texture Panel, plus a Default slot that acts simply as a flat texture.

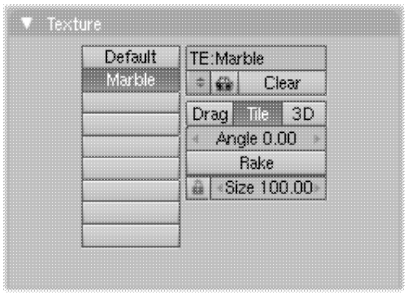
The Brush panel gives you some options applied whatever is the texture you chose:

Curve and Reset
: The Curve button allows you to use the curve control to the right to modify the intensity of the brush from its centre (left part of the curve) towards its borders (right part of the curve). The default curve, which can always be Reset with that button, gives you a smooth brush. When Curve is disabled, you will get the same result as if the curve was constant **1** (a horizontal line at the top).

Space
: Setting this to a non-zero value adds extra space between each application of the brush. The value is measured in pixels (on screen); setting Space to **100** will require the mouse to move **100** pixels between each â← dotâ← applied to the mesh. Note that this is the total distance the brush has travelled, not the current linear distance from the last time the brush was applied.

View
: By default, we saw that vertices are moved along their normals. This option allows you to move them more or less towards the view (**0** has no effect, **5** is half effect, **10** is full effect). Only available with the Draw brush.

Anchored (â § ShiftA)
: This toggle option modifies how the brush behaves. When enabled, the centre of the brush is *anchored* to its initial location (where it was when you clicked). Moving the mouse just enlarges/reduces the brush size, and rotates it. For example, this allows you to apply uniformly a 3D pattern (*via* the texture) to a whole area of your model.



The Texture panel.
Any texture type can be loaded into one of the Sculpt Mode texture slots. Once a texture is associated with a slot, additional options will appear that affect how the texture controls the brush.

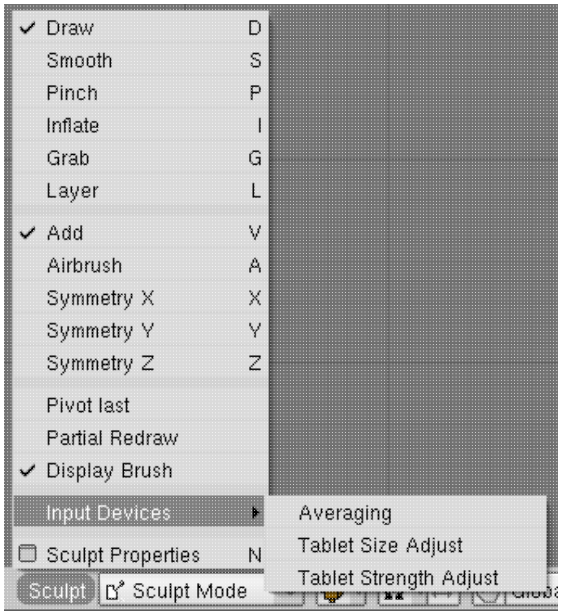
Drag, Tile and 3D
: These three options control how the texture is mapped onto the brush. If Drag is enabled, the texture follows the mouse, so it appears that the texture is being dragged across the model. The Tile option tiles the texture across the screen, so moving the brush appears to move separately from the texture. The Tile option is most useful with tileable images, rather than procedural textures. Lastly, the 3D allows the brush to take full advantage of procedural textures. This mode uses vertex coordinates rather than the brush location to determine what area of the texture to use.

Angle
: This is the rotation angle of the texture brush. It can be changed interactively via CtrlF in the 3D view. While in the interactive rotation you can enter a value numerically as well.

Rake
: When enabled, the brush will be automatically rotated to follow the movement of the tool. Not available with 3D textures.

Size
: This setting allows you to modify the scaling factor of the texture. Not available for Drag textures.

Sculpt Menu



The Sculpt menu.
The Sculpt menu offers several new controls in addition to the tools already discussed.

Pivot Last
: Sets the rotation centre for rotating the scene to the last location the brush was used.

Partial Redraw
: This uses a special graphics optimisation that only redraws where the mouse has been - it can speed up drawing on some graphics cards but slow it down on others. Primarily is only useful needed for dense mesh (greater than 100,000 polys).

Display Brush
: Controls whether the brush circle is drawn.

Input Devices
: Here you can select the behaviour of the used input devices.
 - ◊ Averaging - This option uses an average direction of movement for the number of pixels specificed and then interpolates the stroke along a linear path of that number of pixels. This can be useful for dense meshes but the speed hit can be such that it may be faster to leave it at 1 (the default).

 The following tow settings are specific to graphic tablets:
 - ◊ Tablet Size Adjust - Sets to what extent tablet pressure affects brush size.
 - ◊ Tablet Strength Adjust - Sets to what extent tablet pressure affects brush strength.

Keyboard Shortcuts

Action	Shortcut
Hide mesh outside selection	Ctrlâ § ShiftLMB
Hide mesh inside selection	Ctrlâ § ShiftRMB
Show entire mesh	AltH
Toggle airbrush	A
Interactively set brush size	F
Interactively set brush strength	â § ShiftF
Interactively rotate brush texture	CtrlF
Toggle brush direction (Add/Sub)	V
Draw brush	D
Smooth brush	S
Pinch brush	P
Inflate brush	I
Grab brush	G
Layer brush	L
Flatten brush	T
X Symmetry	X
Y Symmetry	Y
Z Symmetry	Z
Toggle floating sculpt panel	N
Step up one multires level	Page up
Step down one multires level	Page down

Overview

Retopo (remake topology) is a tool for remaking the topology of a mesh. Itâ← s the opposite of the sculpting tools: instead of reshaping the model but leaving the topology the same, it reforms the topology, but maintains the same shape as the original model. You will not change the geometry of the original mesh in any way, you will be creating a new mesh that is projected upon an existing mesh. There are three ways to use retopo:

- Retopo paint, where you paint lines on the mesh and painted intersections become vertex locations.
- Creating new mesh via standard mesh editing methods (the new vertices will â← stickâ← to the reference objectâ← s surface).
- Projecting an existing mesh onto the object.

In practice one usually uses all three together.

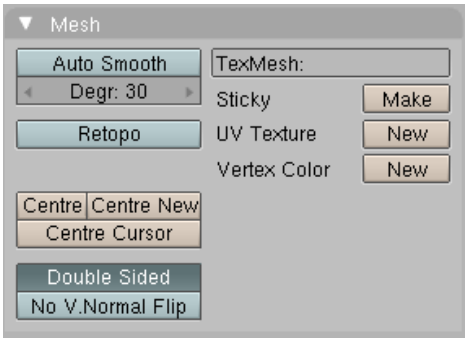
Usage

Retopo is controlled from the Mesh panel in the Editing context (F9).

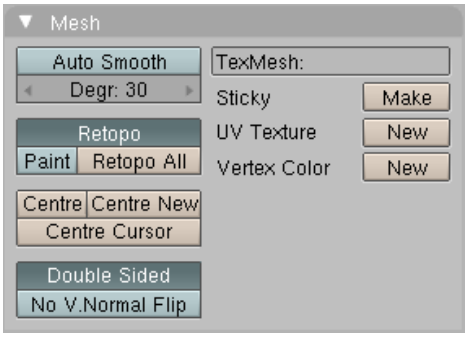
The Retopo button appears when a mesh is selected and Edit mode is active.

When the Retopo toggle button is pressed, any change to a vertex will cause it to be snapped to the surface of *another* model. Note that this effect is view dependent: from the view you are working in, the vertices wonâ← t appear to have moved, because they are falling straight back along the axis until they hit the surface of the model.

Retopo Paint Overview

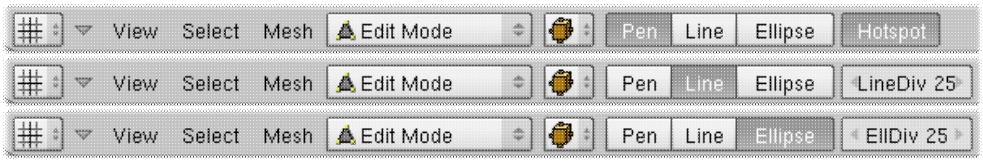


Mesh panel with Retopo disabled.

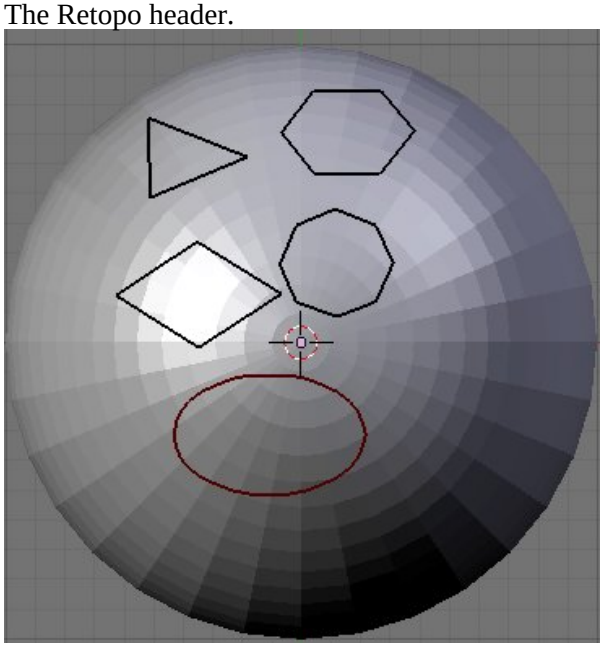


Mesh panel with Retopo enabled.

Once you click Retopo and then Paint, you will see the following options in the header:



The Retopo header.



Drawing ellipses.

Pen

A freehand line that exactly follows the location of the mouse from your start click (LMB) until you release the mouse button.

HotSpot

When enabled, this option allows you to easily connect a new line to one end of another one, by clicking when your mouse pointer is sufficiently close to this end (you see a small grey circle around it). You can toggle this option with H.

Line
: A line segment from your start click (LMB) to where you release the mouse button.

LineDiv
: This setting control the number of sub-segments in a line, the number of points used to project it onto the object. Thus with two points only the endpoints are used, and the rest of the line does not follow the surface but instead passes through the object. And higher levels of LineDiv mean the projection will more accurately fit on the objectâ← s surface.

Ellipse
: Draws a ellipse between a rectangular area with its centre defined by your start click (LMB) and a corner defined by where you release the mouse.

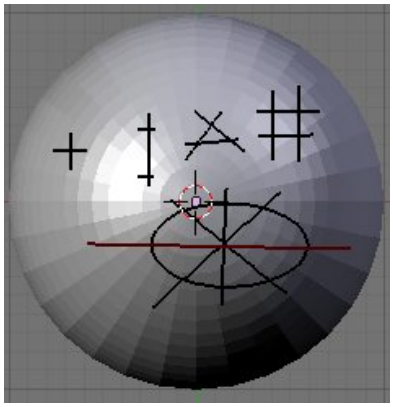
EllDiv
: This setting control the number of sub-segments in the ellipse line. Thus a setting of 4 for EllDiv results in a diamond being draw instead of an ellipse.

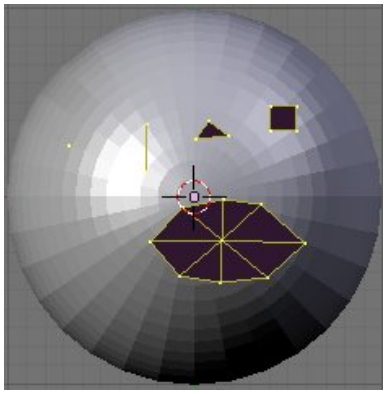
To extend a previous Pen stroke, enable HotSpot, go to an endpoint of the stroke and a circle will appear â← click and drag to continue the stroke.

You can delete previous strokes of any kind by first activating the Pen tool, then going to an endpoint (or for an ellipse along the edge) and when the circle appears click to turn the stroke red, then X or Delete to delete it.

Sometimes the option to extend can be an annoyance if you are trying to paint a new stroke near the endpoint of an existing stroke. In which case you can toggle off the hotspot for end point extension with the H.

After painting your ellipses and lines press â µ Enter to convert them to mesh. At each intersection of lines a vertex is formed. Two intersections on the same line forms an edge also. Three or four edges in a loop form a face.

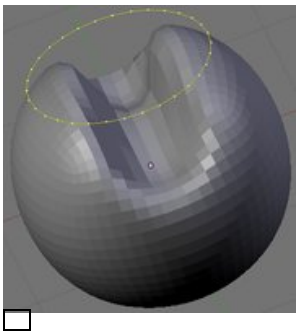

Drawing different shapes (mouse strokes).


Drawing different shapes (results).

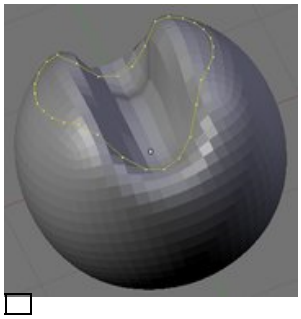
Limitations / Notes

- When you press â µ Enter the mesh faces are derived based on your current view; thus you can only paint one side at a time.
- Lines cannot self intersect
- Can only autofill quads and tris, if an area has more than four surrounding vertices then a face will not be created
- New mesh from paint and previous mesh from paint must be stitched together manually.
- Each stroke is sequential.
- Paint strokes can only be viewed and edited in the 3D view you begin retopo paint in.

Retopo Mesh Overview



Circle before projection.



Circle projected onto carved sphere.

If you use Retopo without Paint you can use standard mesh editing tools to either create a new mesh. Or alternatively project an existing mesh onto the surface you wish to remake the topology of.

To project an existing mesh select the vertices you wish to project in the view you plan to project from then press Retopo All.

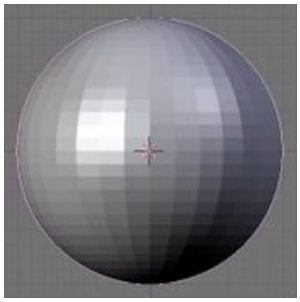
You may wish to hide (H) the vertices you have already projected in order to avoid accidentally re-projecting them from a different view.

Limitations

- In order to avoid affecting mesh on the backside of the object you must hide (H) these vertices (i.e. selection is not â← clippedâ← even if it is not visible).
- To re-topologize part of your current mesh you must either duplicate the entire object in Object mode or duplicate and separate part of the object in Edit Mode.

Tutorial

Step One



Step One.
Since retopo modifies topology, not shape, we must first create the shape we will be using. This shape can be almost any 3D object: a mesh, a NURBS surface, metaballs, or even bones. One useful workflow is to quickly block out a shape in Sculpt mode, then use retopo on that mesh.

For this example, Iâ← ve chosen a simple UVsphere.

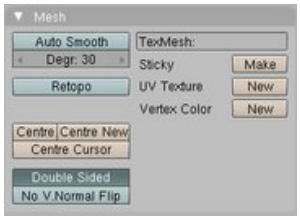
Step Two



Step Two.
In Object mode, add a new mesh. It doesnâ← t matter what you choose. After you create it, press X in *Edit* mode. From the pop-up menu that appears, choose *Erase..Vertices* to erase all vertices. (You don't need those vertices, because you're about to make your own vertices in the next steps.)

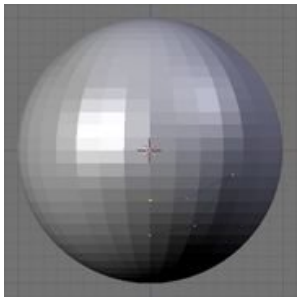
(**Note:** if the popup instead says *"Erase selected object(s),"* you're not in Edit mode. You want to delete the unnecessary *vertices*, but not the *object.*)

Step Three



Step Three.
Go in Edit mode (â Tab) and turn on Retopo. The Retopo toggle is in the Mesh panel in the Editing context (F9). You should also check that you have Viewport Shading set to Solid and that Occlude background geometry is off (thatâ← s the cube icon next to the vertex/edge/face selection buttons, in the 3D view header).

Step Four



Step Four.
Start adding points by clicking CtrlLMB. This is a normal Edit Mode operation, except that if you now rotate your view, you can see that the vertices you just added are on the surface of the UVsphere. You can also use extrude, duplicate, grab, rotate, and scale; all of these operations will continue to snap vertices to the surface of the object behind them.

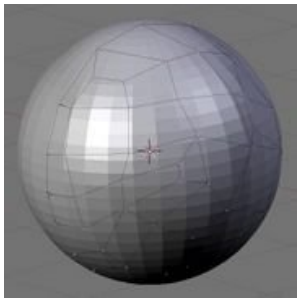
Tips:

To make using Retopo easier, make sure youâ← re taking advantage of Blenderâ← s theme settings. You can use them to increase the size of vertices or to give better contrast to vertices and edges. Note you canâ← t modify default theme â← you should add and edit a new copy, see this page.

You may also find it helpful to turn on the X-ray button in the Object context (F7), Draw panel. This way, your new mesh will always be visible, even when behind another object.

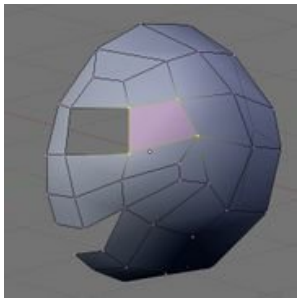
And you may find it useful to have multiple 3D viewports open so that you can see whether the vertex placement is as you desired without needing to rotate your working view around the model.

Step Five



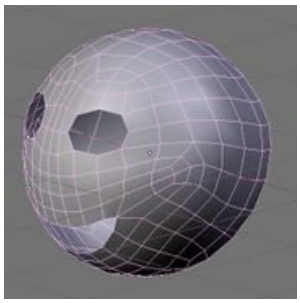
Step Five.
Continue adding points around the rest of the model. Make sure to connect vertices with edges so that you can see the topology you are creating.

Step Six



Step Six.
Once you have all the vertices and edges created, you can turn off Retopo and hide the UVsphere. Still in Edit mode, start selecting groups of edges and use F to add faces. When youâ← re done, you should have a complete surface.

Step Seven



Step Seven.
Add a Mirror modifier and, optionally, a Subsurf one.

Curves

Curves and Surfaces are objects just as meshes are objects except they are expressed in terms of mathematical functions, rather than as series of points.

Blender implements both BÃ©zier curves and Non Uniform Rational B-Splines (NURBS) curves and surfaces. Both are defined in terms of a set of â control verticesâ which define a â control polygonâ , though each follow a different set of mathematical laws. The way the curve and the surface are interpolated might seem similar, at first glance, to Catmull-Clark subdivision surfaces. The curve is *interpolated* while the surface is *attracted*.

When compared to meshes, curves and surfaces have both advantages and disadvantages. Because curves are defined by less data, they produce nice results using less memory at modelling time, whereas the demands increase at rendering time.

Some modelling techniques, such as extruding a profile along a path, are only possible with curves. But the very fine control available on a per-vertex basis on a mesh, is not possible with curves.

There are times when curves and surfaces are more advantageous than meshes, and times when meshes are more useful. If you have read the pages on Basic Mesh Modeling and Advanced Mesh Modeling, and after you read this part, you will be able to choose whether to use meshes or curves.



Logo Thumbnail.
Working with curves in Blender is fairly simple and surprisingly there are very few HotKeys when creating curves. It is what you do with those curves that really makes the difference. A curve by itself is just that, a curve. But a curve applied to another curve can create very complex objects.

When you have finished reading and learning about BÃ©zier and NURBS curves there are several more advanced examples on the application of curves in the tutorials section for modelling complex objects.

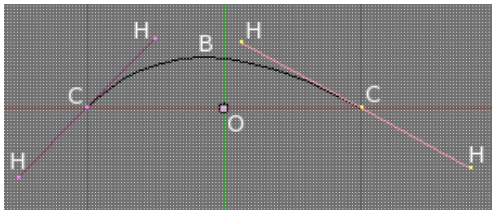
There is a Working example that shows how to create an interesting bird-like logo, (*Logo Thumbnail*). The tutorial covers most aspects of working with BÃ©zier curves including: adding curves, setting up a background image as a template guide and bevelling the final curve.

In addition, the Tutorial section has examples on both Skinning and Curve deform techniques.

BÃ©ziers

BÃ©zier curves are the most commonly used curve for designing letters or logos. They are also widely used in animation, both as paths for objects to move along and as IPO curves to change the properties of objects as a function of time.

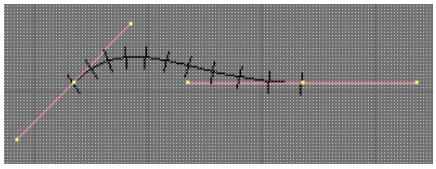
There are three panels designed to assist in working with and modifying curves: Curve and Surface, Curve Tools and Curve Tools1. Each panel has buttons that change the characteristics of curves.



Curve example.
(*Curve example*) is the most basic curve you can create. It consists of two control points or vertices, labelled â← Câ← , the curve â← Bâ← , handles â← Hâ← and an object centre â← Oâ← .

Selecting the control point also selects the handles, and allows you to move the complete vertex. Selecting one or more handles allows you to change the shape of the curve by dragging the handles.

To create a curve use the toolboxâ← s Add â Curve â Bezier Curve menu entry to add a new curve, (*Curve example*). By default the new curve exists only in 2D. For example, if you created the curve in the Top view, the shape of the curve can only be change in the XY Plane. You can apply transforms to the curve but you canâ← t change its shape in 3D.

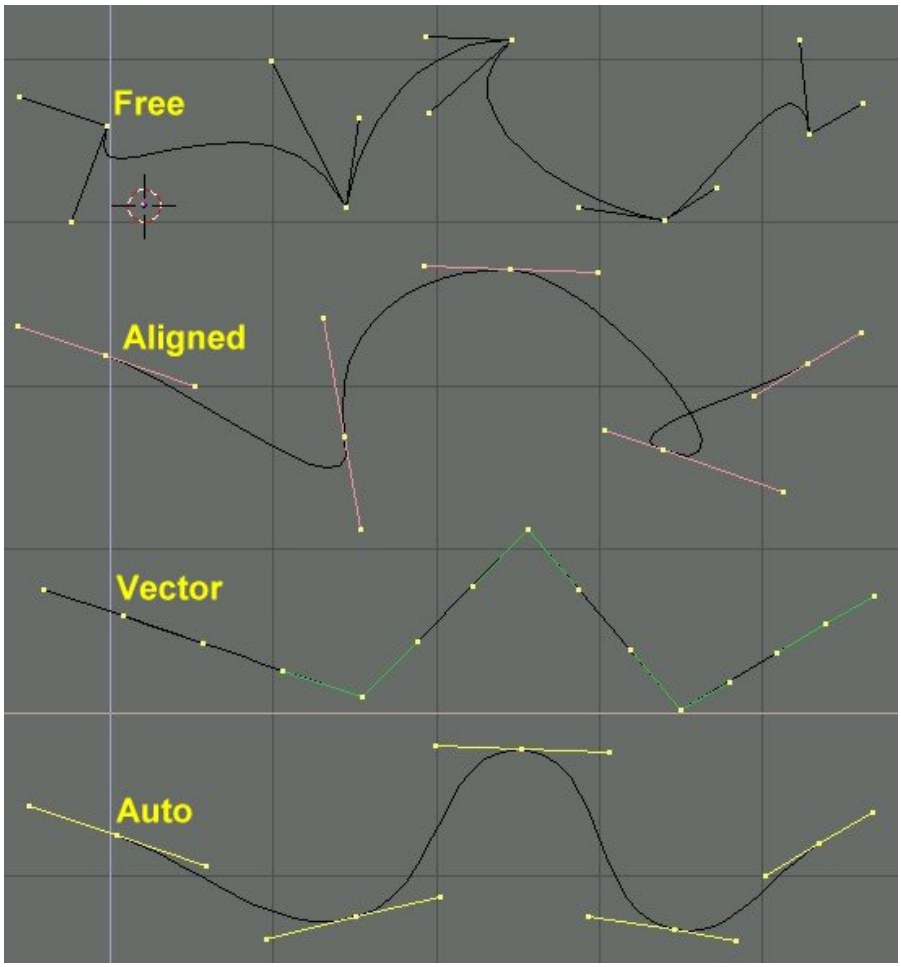


3D Curve - a Path.
To work with the curve in 3D you need to *turn on* the 3D property of the curve using the 3D button in the Curve and Surface panel. You can visually see that a curve is in 3D by noticing the curve has *railroad* tracks or marks. (*3D Curve - a Path*) is a 3D curve and (*Curve example*) is a 2D curve.

A handle is always tangent to the curve. The â← steepnessâ← of the curve is controlled by the handleâ← s length, any â← Hâ← to a â← Câ← . The longer a handle is the steeper the curve (i.e. the more curve wants to *hug* the handle).

There are four types of handles (*Types of Handles for BÃ©zier curves*):

- Free Handle (black). The handles are independent of each other. To convert to Free handles use H. H also toggles between Free and Aligned.
- Aligned Handle (purple). These handles always lie in a straight line. Hotkey: H (toggles between Free and Aligned).
- Vector Handle (green). Both parts of a handle always point to the previous handle or the next handle. Hotkey: V.
- Auto Handle (yellow). This handle has a completely automatic length and direction, set by Blender to ensure the smoothest result. Hotkey: â § ShiftH.



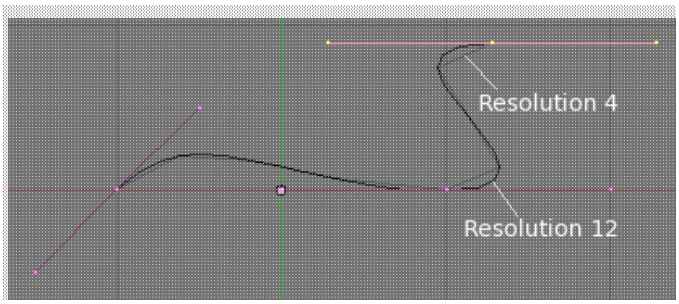
Types of Handles for BÃ©zier curves.

Handles can be *grabbed*, *rotated* and *scaled* exactly as ordinary vertices in a mesh would. As soon as the *handles* are moved, the handle type is modified automatically:

- Auto Handles becomes Aligned;
- Vector Handles becomes Free.

Curve resolution

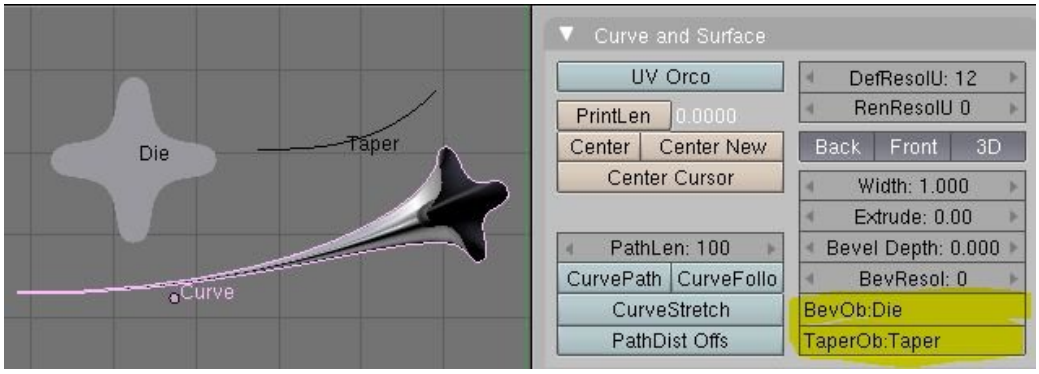
Although the BÃ©zier curve is a continuous mathematical object it must nevertheless be represented in discrete form (set of small segments) from a rendering point of view. This is done by setting a *resolution* property, which defines the number of points which are computed between every pair of control points.



Resolution example.

A separate *resolution* can be set for each BÃ©zier curve by adjusting the DefResolU field. The default is **6**. (*Resolution example*) is an example of the same curve, superimposed, with the aid of Gimp, showing two different resolution settings. The lighter shaded curve has a low resolution of **4**; the curve begins to look linear. The darker curve has a resolution of **12** and is very smooth. Note that high resolutions may look nicer but they can slow down interactive rendering if there is a large number of curves (et/or if these ones have a large amount of control points).

Bevel and Taper Objects



3D Curve modified by Bevel and Taper curves.
A Bevel object, applied to a Curve object, forms a skin for the curve. Where the curve is the path or length of a pipe, the Bevel Object, extruded along that path, defines the outside shape, like the outside of a cord, or hose pipe. Normally the Bevel is a simple round circle, and thus makes the curve into a pipe or soda can. The Bevel shape must be two-dimensional, and it can be rectangular for simulating wrought iron or flat steel, oval (with a crease) for a power cord, star-shaped for a shooting star illustration; anything that can be physically formed by extrusion (extruded).

A Taper object is an open curve with control points above its object centre. When applied to a Bevelled Curve, it changes the diameter of the Bevel along the length of the curve, like a python just having eaten a rat, or like a hose bulging up under pressure, or a vine growing.

For adjusting proper size of Bevel effect for individual curveâ← s segment use Set Radius option accessible through W-4. Default value is **1.0**.

Caution: no Bevel Effect if bevel Radius parameter set to **0.0**.

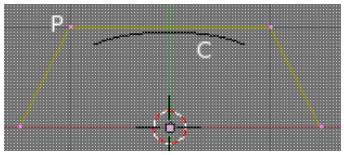
NURBS

NURBS curves are defined as rational polynomials and are more general, strictly speaking, than conventional B-Splines and BÃ©zier curves inasmuch as they are able to exactly follow any contour. For example a BÃ©zier circle is a polynomial *approximation* of a circle, and this approximation is noticeable, whereas a NURBS circle is *exactly* a circle.

NURBS curves require a little bit more understanding of the underlying components that make up a NURBS curve in order to get full use of them. They have a large set of variables, which allow you to create mathematically pure forms. However, working with them requires a little more discussion on the various parts of a NURBS curve.

Uniform-Endpoints

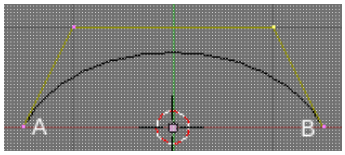
We start with *Knots*. NURBS curves have a *knot* vector, a row of numbers that specifies the parametric definition of the curve (i.e. they describe the range of influence for each of the control-points). Remember the control-points from BÃ©zier curves, NURBS have them too and each control-point affects some part of the curve along that range. The control-points appear as purple vertices.



Default Uniform curve.
(*Default Uniform curve*) is the default NURBS curve created using the NURBS Curve menu item from the toolboxâ← s Add menu and is an example of a Uniform curve. The curve itself is drawn in black, labelled â← Câ← and the control-points are drawn in purple; one out of the 4 is labelled â← Pâ← .

You canâ← t manipulate the *Knot* vector directly but you can configure it using two presets: Uniform and Endpoint.

The Uniform button produces a uniform division for closed curves, but when used with open curves you will get â← freeâ← ends, which are difficult to locate precisely.

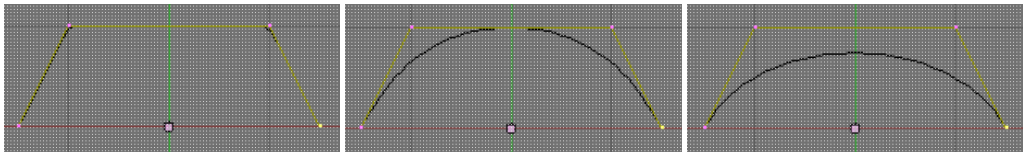


Endpoint curve.
The Endpoint button sets the *Knot* vector in such a way that the first and last vertices are always part of the curve, which makes them much easier to place. (*Endpoint curve*) is an example of applying the Endpoint button to the (*Default Uniform curve*). You can see that the curve has now been pulled to the *end* control-points labelled â← Aâ← and â← Bâ← .

Order

The Order field is the â← depthâ← or *degree* of the curve (i.e. you are specifying how much the control-points are taken into account for calculating the curve shape).

Order **1** is a point and is not an available depth setting, Order **2** is linear (*Order 2 curve*), Order **3** is quadratic (*Order 3 curve*), (*Order 4 curve*) is cubic, and so on. The valid range is **2** to **6**. Notice that as the Order rises the curve moves away from the control-points.



Order 2 curve. Order 3 curve. Order 4 curve.

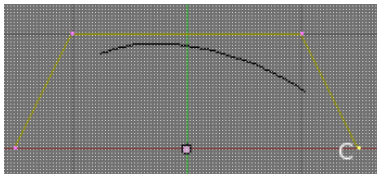
If your curve has **6** or more control-points the Order can not be set higher than **6**. **6** is the highest Order allowable. If you have less than **6** control-points then the highest Order is limited by the number of control-points. For example, if your curve has **5** control-points then the highest Order allowable is **5**.

Always use an order of **5**, if possible, for curve paths because it behaves fluidly under all circumstances, without producing irritating discontinuities in the movement. For example, if you have a cube assigned to travel along a NURBS path with an Order of say **2** then the cube will appear to move roughly (or jerky) along the path.

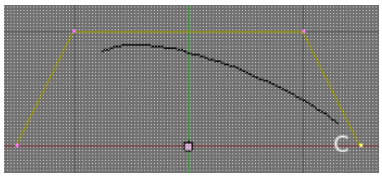
Math Note
Mathematically speaking the Order is the order of both the Numerator and the Denominator of the rational polynomial defining the NURBS curve.

Weight

NURBS curves have a Weight assigned to each control-point that controls how much each â← pullsâ← on the curve. Think of it as if each control-point has an arm that reaches out and grabs hold of the curve and tries to pull on it. The larger the Weight the more the control-point pulls on the curve, see (*Weight of 5*) and (*Weight of 20*). The valid range of Weight settings are **0.1** to **100.0**.

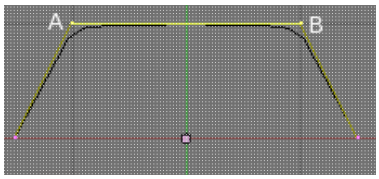


Weight of 5.



Weight of 20.

The larger Weight of 20 pulls the curve towards the control-point labelled â← Câ← . Each control-point can have a different Weight setting. As the Weight for a control-point increases the curve will hug the control-points closer. If the Weights are large enough the curve will almost follow the control-points, see (*Weight of 100*).



Weight of 100.

The control-points can effectively compete with each other. For example, the control-point with the largest Weight will pull the curve towards it and away from the others. If all the control-points have the same Weight then the Weight is effectively cancelled, as if none had Weights.

In (*Weight of 100*) the top two control-points have their Weight set at **100.0**, labelled â← Aâ← and â← Bâ← . The opposite control-points have their Weight at **1.0**. You can see that the curve is â← pulledâ← toward control-points â← Aâ← and â← Bâ← . And at such a high Weight the curve almost follows the control-points.

In the Curve Tools panel there are four preset Weights that provide typical Weight settings for certain kinds of control-point arrangements. Some generate Weight settings that are used for control-points that form circles.

To see the Weight value of a control-point, select it, open the Transform Properties panel using N and look at the Vertex W field. The Weight field doesnâ← t show the Weight!

Resolution

As with BÃ©ziers curves NURBS curvesâ← Resolution can be controlled from the Curve Tools panel.

Opening, Closing, Deleting, Joining, Bevel, Taper

As with BÃ©ziers curves, Opening, Closing, Deleting and Joining NURBS curves is performed using the same Hotkeys and same Curve Tools, with the same rules applying, see the BÃ©ziers curves section above, as well as the next page.

Add new segment

Mode: Edit mode

Hotkey: CtrlLMB or E

Menu: Curve â Extrude

Description

Once a curve is created you can add new segments, either by extruding it, or placing new handles with CtrlLMB clicks. Each new segment is added to the end of the curve. A new segment will only be added if a single vertex, or handle, at one end of the curve is selected. If two or more vertices are selected nothing is added.

Opening and Closing a Curve

Mode: Edit mode

Hotkey: C

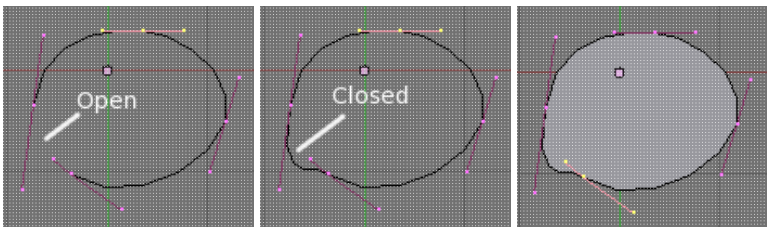
Menu: Curve â Toggle Cyclic

Description

This toggles between an open curve and closed curve. The shape of the closing segment is based on the handles. The only time a handle is adjusted after closing is if the handle is an Auto handle. (*Open curve*) and (*Closed curve*) is the same curve open and closed.

This action only works on the original starting control-point or the last control-point added. Deleting a segment(s) dosenâ← t change how the action applies; it still operates only on the starting and last control-points. This means that C may actually join two curves instead of closing a single curve.

Examples



Open curve. Closed curve. Closed curve (Solid).

If the curve is closed the curve is automatically considered a surface with an area. This means it is rendered as a solid (*Closed curve (Solid)*) and it is renderable using F12.

Deleting a Segment(s)

Mode: Edit mode

Hotkey: X

Menu: Curve â Delete...

Description

A segment is defined implicitly by its two adjacent control-points. You canâ← t explicitly select a segment; you must select this two adjacent control-points. Once the control points are selected you can use the Erase menu and select the Segment menu item.

Note, that unlike with *meshes*, deleting one or more handles (Selected item) do not â← cutâ← the curve : the adjacent segments are just joined.

Joining two curves

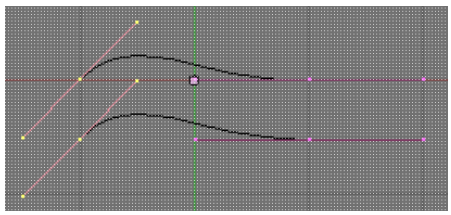
Mode: Edit mode

Hotkey: F

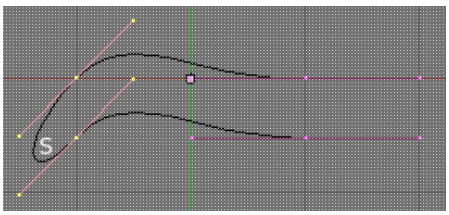
Menu: Curve â Make Segment

Description

Joining curves is really the act of making a segment between the two curves. To join two separated use one control point from each curve. The two curves are joined by a segment to become a single curve. (*One curve joined*) is the result of joining (*Two curves*). The segment, labelled â← Sâ← , is the new segment joining the two curves. We use F for the hotkey because it is similar to making a new **F**ace in a mesh.



Two curves.



One curve joined.

You can not close a curve by joining the curves; you must Close the curve. You will get the error â← `Can't make segment`â← when you attempt to join using the starting and last control-point. For example, in (*One curve joined*) you must use Close to close the curve.

Curve Deform

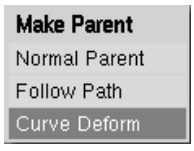
Curve Deform provides a simple but efficient method of defining a deformation on a mesh. By parenting a mesh object to a curve, you can deform the mesh up or down the curve by moving the mesh along, or orthogonal to, the dominant axis.

The Curve Deform works on a (global) dominant axis, X, Y, or Z. This means that when you move your mesh in the dominant direction, the mesh will traverse along the curve. Moving the mesh in an orthogonal direction will move the mesh object closer or further away from the curve. The default settings in Blender map the Y axis to the dominant axis. When you move the object beyond the curve endings the object will continue to deform based on the direction vector of the curve endings.

A Tip:

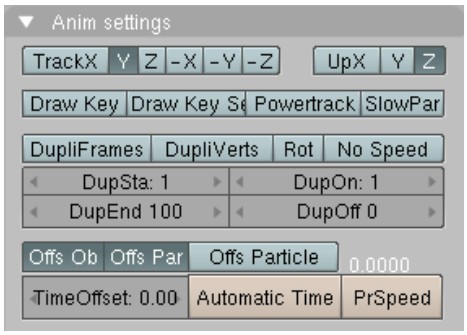
Try to position your object over the curve immediately after you have added it, before adding the curve deform. This gives the best control over how the deformation works.

Interface

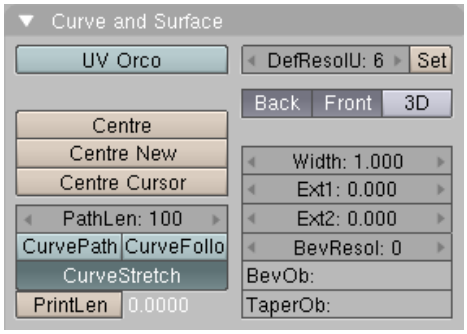


Make Parent menu.
When parenting a mesh to a curve (CtrlP), you will be presented with a menu (*Make Parent menu*).

By selecting Curve Deform, you enable the curve deform function on the mesh object.



Anim settings panel.
The dominant axis setting is set on the mesh object. By default the dominant axis in Blender is Y. This can be changed by selecting one of the Track X, Y or Z buttons in the Anim Panel, (*Anim settings panel*), in Object context (F7).



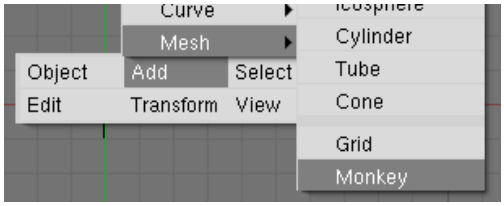
Curve and Surface panel.
Cyclic curves work as expected where the object deformations traverse along the path in cycles.

The CurveStretch curve option allows you to let the mesh object stretch, or squeeze, over the entire curve. This option is in Editing Context (F9), for the curve. See (*Curve and Surface panel*).

This is also a modifier
The Curve Deform relationship is now also a modifier, called Curve. The Curve modifier function acts the same as its counterpart, except that when the modifier is used, the â← dominant axisâ← is set inside its properties â← and the Track X/Y/Z buttons have no more effect on it.

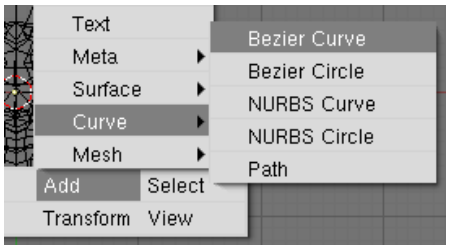
Example

Letâ ← s make a simple example:



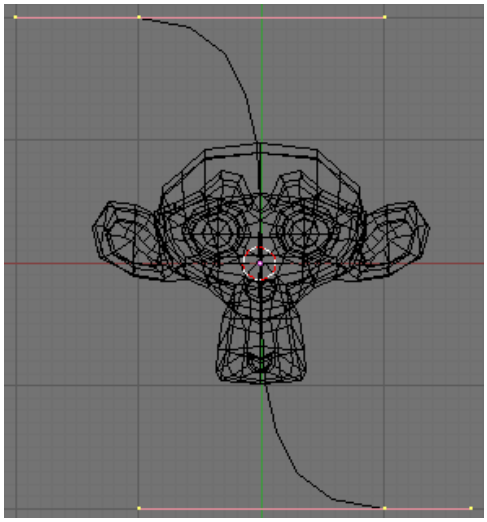
Add a Monkey!

- Remove default cube object from scene and add a Monkey (â § ShiftA â Add â Mesh â Monkey, *Add a Monkey!*)!
- Now press â Tab to exit Edit Mode.



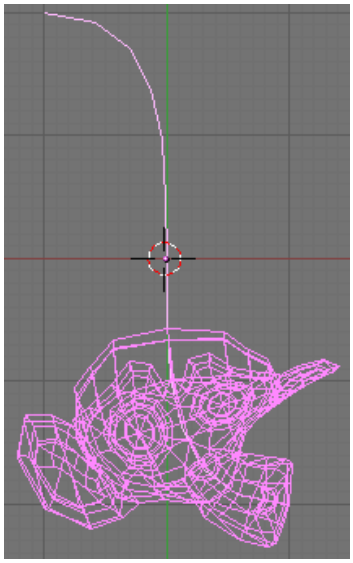
Add a Curve.

- Now add a curve (â § ShiftA â Add â Curve â Bezier Curve, *Add a Curve*).



Edit Curve.

- While in Edit mode, move the control points of the curve as shown in (*Edit Curve*), then exit Edit mode (â Tab).



Monkey on a Curve.

- Select the Monkey (RMB), and then shift select the curve (â § ShiftRMB).
- Press CtrlP to open up the Make Parent menu.
- Select Curve Deform (*Make Parent menu*).
- The Monkey should be positioned on the curve, as in (*Monkey on a Curve*).
- Now if you select the Monkey (RMB), and move it (G), in the Y-direction (the dominant axis by default), the monkey will deform nicely along the curve.

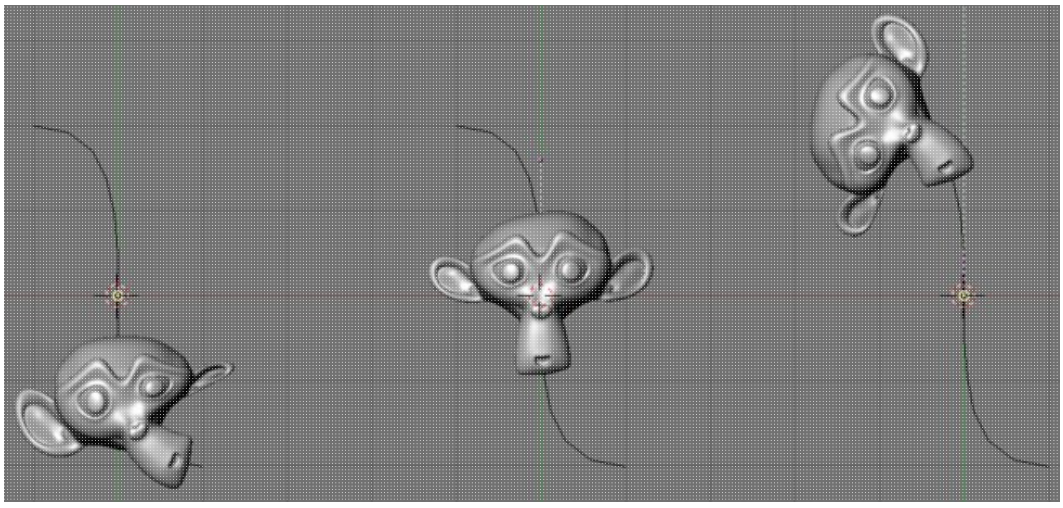
A Tip:

If you press MMB while moving the Monkey you will constrain the movement to one axis only.

- In (*Monkey deformations*), you can see the Monkey at different positions along the curve. To get a cleaner view over the deformation I have activated SubSurf with Subdiv to **2**, and Set Smooth on the Monkey mesh (F9 to get Editing context).

A Tip:

Moving the Monkey in directions other than the dominant axis will create some odd deformations. Sometimes this is what you want to achieve, so youâ← ll need to experiment and try it out!



Monkey deformations.

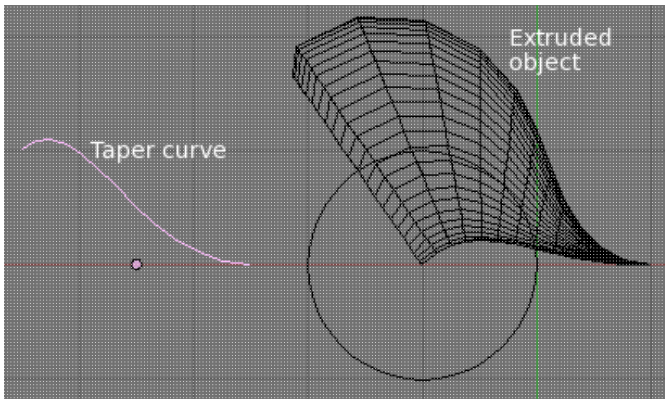
Curve Taper

Taper is a scaling effect where the scaling is controlled by a curve. The tool is used on bevelled or extruded objects. As with Extruding, Tapering is â activatedâ by entering the name of a curve in the TaperOb field on the Curve and Surface panel.

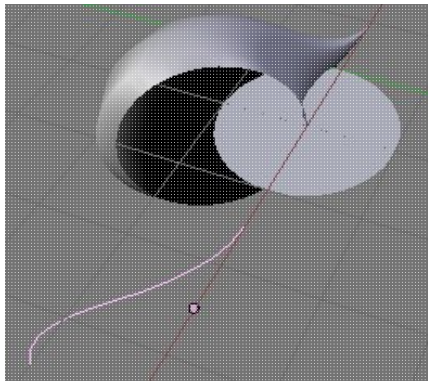
The taper curve defines the width of the extrusion along the curve on the â Bevelled Objectâ that is defined in the BevOb field. The taper curve is typically horizontal, where the height (local Y coordinates) denotes the *scale* being applied. If a point on the taper curve goes above **0**, on the local Y axis, enlargement occurs. If a point on the taper curve goes below **0** then shrinking occurs. Think of the taper curve as changing the volume of the bevelled object.

For example we can taper a simple curve circle extruded object using a taper curve. Add a curve, exit Edit mode, add another one (a closed one, like a circle), call it â `BevelCurve`â , and enter its name in the BevOb field of the first curve (Editing context F9, Curve and Surface panel). We now have a pipe. Add a third curve while in Object mode and call it â `TaperCurve`â . Adjust the left control-point by raising it up about 5 units.

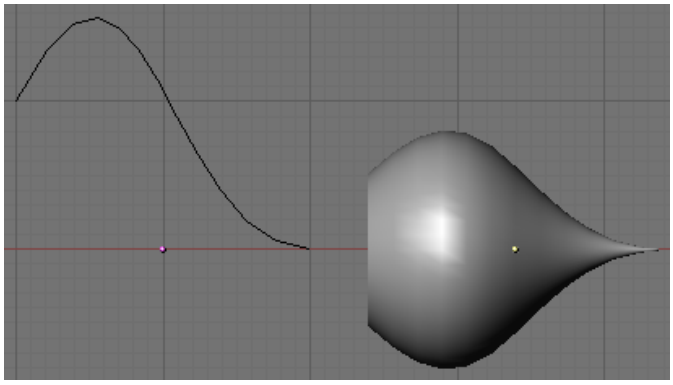
Now return to the Editing context, and edit the first curveâ s TaperOb field in Curve and Surface panel to reference the new taper curve which we called â `TaperCurve`â . When you hit enter the taper curve is applied immediately with the results shown in (*Taper extruded curve*).



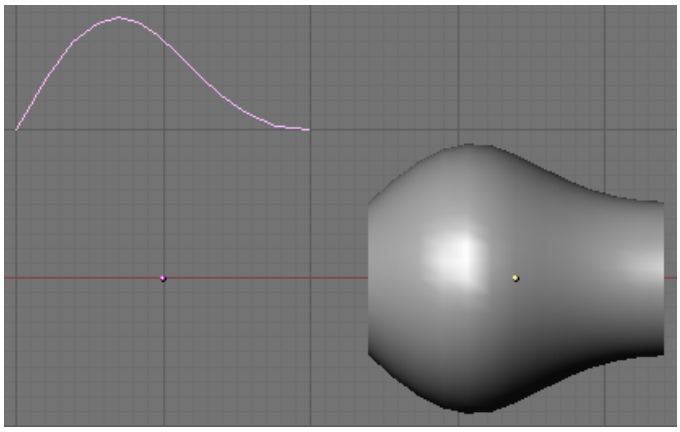
Taper extruded curve.


Taper solid mode.

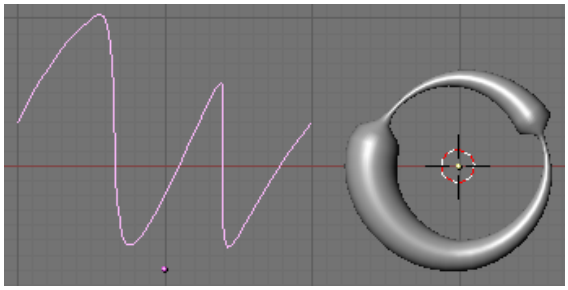
You can see the **Taper curve** being applied to the **Extruded object**. Notice how the pipeâ s volume shrinks to nothing as the taper curve goes from left to right. If the taper curve went below the local Y axis the pipeâ s inside would become the outside which would lead to rendering artifacts. Of course as an artist that may be what you are looking for!



Taper example 1.
In (*Taper example 1*) you can clearly see the effect the left taper curve has on the right curve object. Here the left taper curve is closer to the object centre and that results in a smaller curve object to the right.



Taper example 2.
In (*Taper example 2*) a control point in the taper curve to the left is moved away from the centre and that gives a wider result to the curve object on the right.



Taper example 3.
In (*Taper example 3*), we see the use of a more irregular taper curve applied to a curve circle.

Important rules to consider

- Only the first curve in a TaperOb is evaluated even if you have several separated segments.
- The scaling starts at the first control-point on the left and moves along the curve to the last control-point on the right.
- Negative scaling, (negative local Y on the Taper Curve) is possible as well. However, rendering artifacts may appear.
- It scales the width of normal extrusions based on evaluating the taper curve, which means sharp corners on the taper curve will not be easily visible. Youâ← ll have to heavily level up the resolution

(DefResolU) of the base curve.

- With Cyclic curves (those curves that connect to form a solid object), the Taper curve in TaperOb acts along the whole curve (perimeter of the object), not just the length of the object, and varies the extrusion depth. In these cases, you want the relative height of the TaperOb Taper curve at both ends to be the same, so that the cyclic point (the place where the endpoint of the curve connects to the beginning) is a smooth transition.

Surfaces

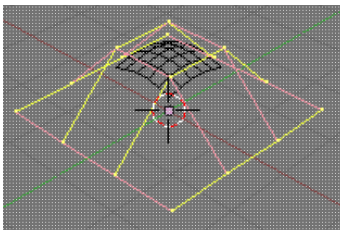
Surfaces are actually an extension of NURBS curves but are still a unique object unto themselves. Whereas a curve produces only one-dimensional interpolation, Surfaces have a second extra dimension of interpolation. The first dimension is U, as for curves, and the second dimension is V.

You may ask yourself â← but the surface appears to be 3D, why is it only 2D?â← . In order to be 3D the object needs to have â← Volumeâ← and a surface doesnâ← t have a volume, it is infinitely thin. If it has a volume the surface would have a thickness. Even though the surface appears to extend in 3D, it has no volume, and hence it only has two interpolation coordinates, U and V. U is the Yellow grid lines and V is the pink grid lines in (*Surface*).

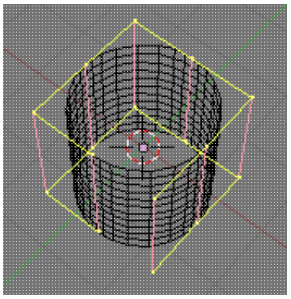
Many of the concepts from NURBS curves carry directly over to NURBS Surfaces, such as control-points, Order, Weight, Resolution, etc.

For NURBS Surfaces the control-points form a grid and is sometimes called a â← Cageâ← . The grid performs exactly like the control-points of a NURBS curve; they control the boundary of the surface.

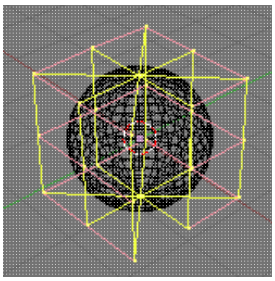
To help get started in creating surfaces there are four preset NURBS Surfaces: (*Surface*), (*Tube*), (*Sphere*) and (*Donut*).



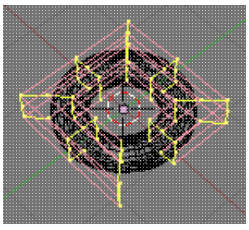
Surface.



Tube.



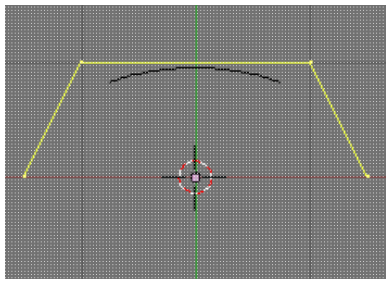
Sphere.



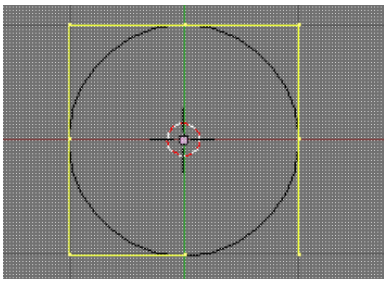
Donut.

Each preset is accessible from the Add â Surface sub-menu of the Toolbox menu where each is designed as a starting point for creating more complex surfaces, of which the most common starting surface is NURBS Surface.

There are also two preset NURBS Surface Curves: (*Curve*) and (*Circle*).



Curve.



Circle.

Although they visually look like NURBS curves they are *not*. Blender internally treats NURBS Surface Curves and NURBS Curves completely differently. There are several attributes that separate them but the most important is that a NURBS Curve has a single interpolation axis and a NURBS Surface Curve has two interpolation axes.

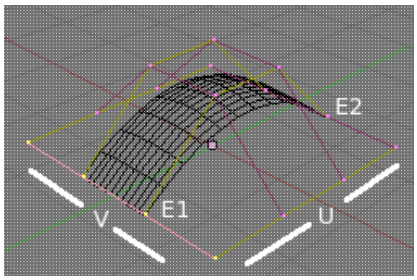
Visually you can tell which is which by entering Edit mode and looking at the 3D windowâ← s header; either the header shows â← Surfaceâ← or â← Curveâ← as one of the menu choices. Also, you can Extrude a NURBS Surface Curve to create a surface but you canâ← t with a NURBS Curve.

Use Surfaces to create and revise fluid curved surfaces. Surfaces can be cyclical in both directions, allowing you to easily create a Donut shape, and they can be drawn as â← solidsâ← in Edit mode. This makes working with surfaces quite easy.

Note
Currently Blender has a basic tool set for Surfaces, with limited ability to create holes and melt surfaces. Future versions will contain increased functionality.

Uniform-Endpoints

Just like with NURBS curves, NURBS Surfaces have a knot vector and the configuration of the knot values are controlled by the Uniform and Endpoint buttons. Each interpolation axis can be independently set to either Uniform or Endpoint.

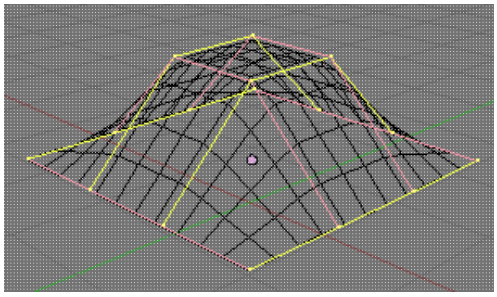


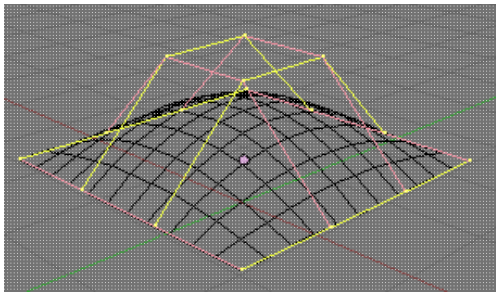
Endpoint U.
In (*Endpoint U*), the U interpolation axis is labelled as â← Uâ← and the V interpolation axis is labelled as â← Vâ← . The Uâ← s interpolation axis has been set to Endpoint and as such the surface now extends to the outer edges from â← E1â← to â← E2â← along the U interpolation axis.

To cause the surface to extend to all edges you would set the Vâ← s axis to Endpoint as well.

Order

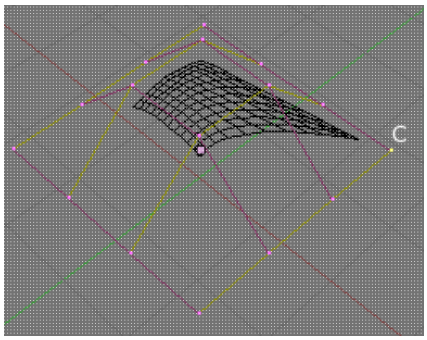
As with NURBS Curves, Order specifies how much the control-points are taken into account for calculating the curve of the surface shape. For high Orders, (*Order 4 surface*), the surface pulls away from the control-points creating a smoother surface â← assuming that the Resolution is high enough. For low Orders, (*Order 2 surface*), the surface follows the control-points creating a surface that tends to follow the grid cage.


Order 2 surface.


Order 4 surface.

For illustration purposes, in both (*Order 4 surface*) and (*Order 2 surface*), the knot vectors were set to Endpoint causing the surface to extend to all edges.

Weight



Surface Weight 100.
Again, as with NURBS Curves, Weight specifies how much each control-point â← pullsâ← on the curve.

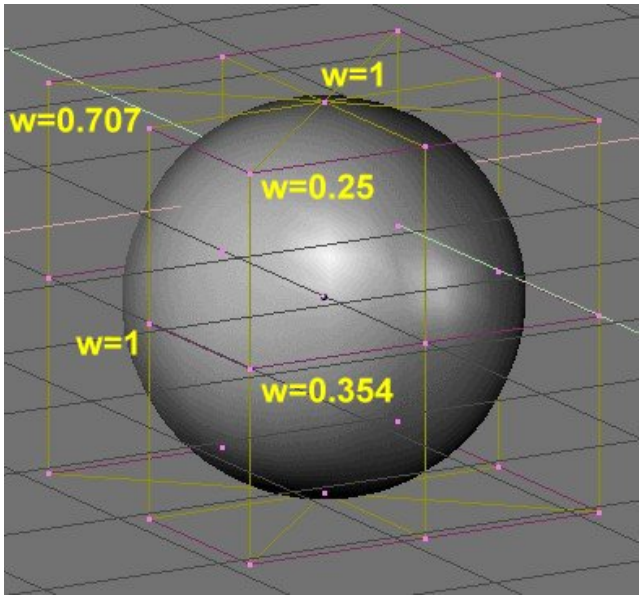
In (*Surface Weight 100*), a single control-point, labelled â← Câ← , has had its Weight set to 100.0 while all others are at their default of 1.0. As you can see that control-point *pulls* the surface towards it.

If all the control-points have the same Weight then each effectively cancels each other out. It is the difference in the Weights that cause the surface to move towards or away from a control-point.

The Weight of any particular control-point is visible in the Transform Properties panel which is accessed using the N.

See NURBS Curves Weight for further details.

Preset Weights



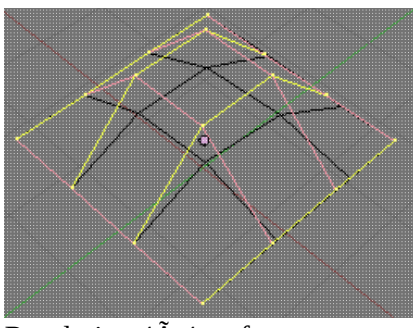
A sphere surface.
NURBS can create pure shapes such as circles, cylinders, and spheres (note that a BÃ©zier circle is not a pure circle). To create pure circles, globes, or cylinders, you must set to specific values the weights of the control-points. This is not intuitive, and you should read more on NURBS before trying this.

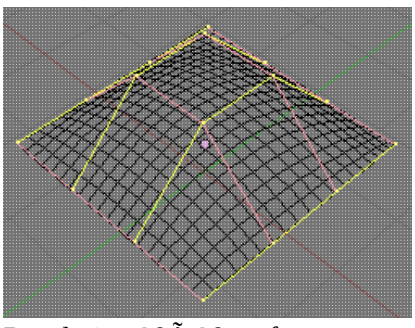
Basically, to produce a circular arc from a curve with three control-points, the end points must have a unitary weight, while the weight of the central control point must be equal to one-half the cosine of half the angle between the segments joining the points. Lets take an example: if you have three control-points creating a right angle, to get a quadrant, you would set the centre pointâ← s weight to `cos(90/2)/2 = cos(45)/2 = sqrt(2)/4 â ↑ 0.354`. This is the same principle for (*A sphere surface*) â← youâ← ll note that the four different weights needed for creating a sphere (**1.0**, **0.707** = `sqrt(0.5)`, **0.354** = `sqrt(2)/4`, and **0.25**) are the four presets available in the Curve Tools panelâ←¦

Resolution

Just like NURBS Curves, Resolution controls the detail of the surface. The higher the Resolution the more detailed and smoother the surface is. The lower the Resolution the rougher the surface.

(*Resolution 4Ã 4 surface*) is an example of a surface resolution of 4 for both `U` and `V`. (*Resolution 20Ã 20 surface*) is an example of a surface resolution of 20 for both `U` and `V`.


Resolution 4Ã 4 surface.


Resolution 20Ã 20 surface.

For illustration purposes the knot vectors where set to Endpoint causing the surface to extend to all edges.

Adding or Extruding

Mode: Edit mode

Hotkey: E

Menu: Surface â Extrude

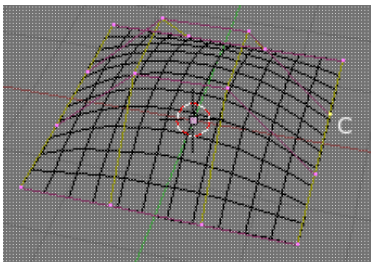
Description

Once the tool is activated the extrusion happens immediately and you are placed into Grab mode, ready to drag the new extruded surface to its destination.

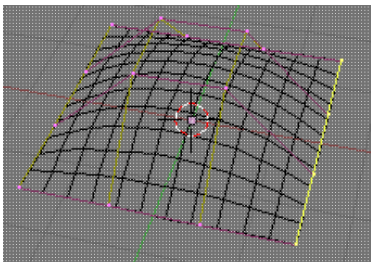
Examples

Images (*Selecting control-point*) to (*Complete*) show a typical extrusion along the side of a surface.

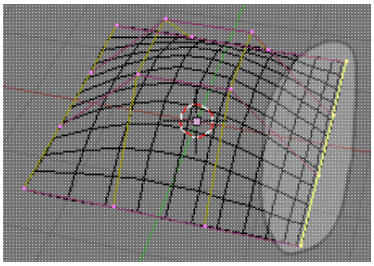
In (*Selecting control-point*) and (*Shift-R*), a row of control-points were highlighted by selecting a single control-point, labelled â← Câ← , and then using the handy row select tool (â § ShiftR) to select the rest of the control-points.

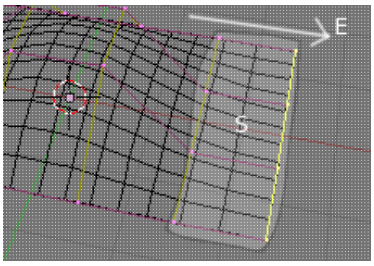


Selecting control-point.


Shift-R.

The edge is then extruded using the E as shown in (*Extruding*). Notice how the mesh has bunched up next to the highlighted edge; the area in question is highlighted in a light-grey circular area. That is because the *new* extruded surface section is bunched up there as well.


Extruding.



Complete.

By moving the new section away from the area the surface begins to â← unbunchâ← , as shown in (*Complete*). The direction of movement is marked with a white arrow, labelled â← Eâ← , and the new section is labelled â← Sâ← .

You can continue this process of extruding â← or adding â← new surface sections until you have reached the final shape for your model.

Cycling (Opening and Closing)

Mode: Edit mode

Hotkey: C

Menu: Surface â Toggle Cyclic

Description

Cycling a surface is similar to Opening and Closing a NURBS curve except that a surface has an inside and outside surface.

To cycle a surface use C and choose either cyclic U or cyclic V from the Toggle menu. The surfaceâ← s outer edges will join together to form a â← closedâ← surface.

Attempting to cycle a non-outer edge will result in nothing happening.

Deleteing/Erasing surfaces

Mode: Edit mode

Hotkey: X

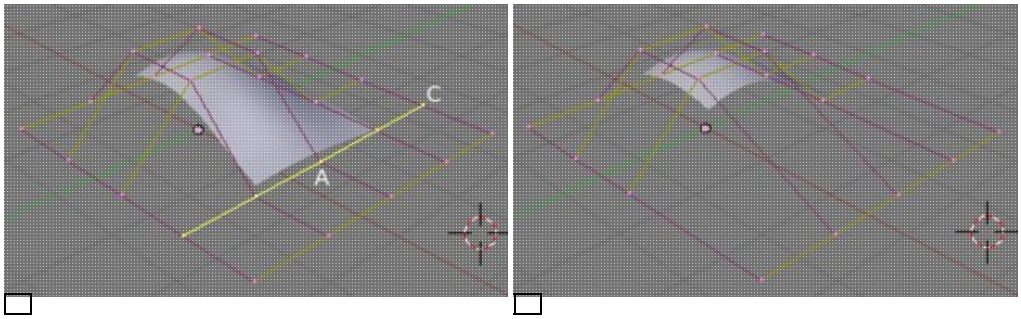
Menu: Curve â Delete

Description

Deleting requires that all control-points along an interpolation axis are highlighted.

Hints

A handy Hotkey (â § ShiftR) is provided that makes it easier to select all the control-points along an axis. Just highlight a control-point(s) and use â § ShiftR to toggle between the two interpolation axes that intersect the last control-point selected.



Before. After.

In (*Before*) a row of control-points have been selected by initially selecting the control-point labelled â← Aâ← and using â § ShiftR to select the remaining control-points. Then, using the Erase menu (X), the *selected* row of control-points is erased resulting in (*After*).

Joining or Merging two surfaces

Mode: Edit mode

Hotkey: F

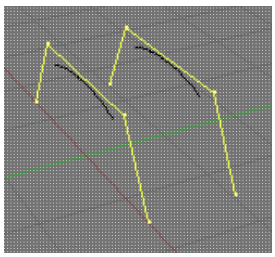
Menu: Surface â Make Segment

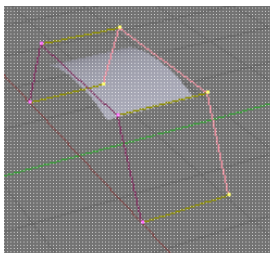
Description

Just like NURBS Curves, Joining requires that a single edge, a row of control-points, from two separate surfaces are selected. This means that the surfaces must be part of the same object. For example, you canâ← t join two surfaces while in Object mode. Joining can only take place while in Edit mode which requires that both surfaces be part of the same object.

Examples

(*Joining ready*) is an example of two NURBS Surface curves, ***not*** NURBS curves, in Edit mode ready to be joined. (*Joining complete*) is the result of joining the two curves.


Joining ready.


Joining complete.

Hints

If not enough surfaces are selected then you will get an error message stating â← `Too few selections to mergeâ←` .

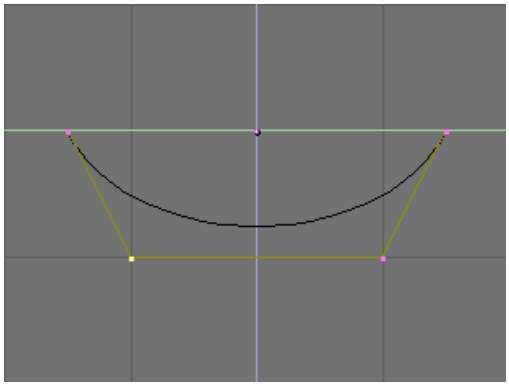
Most of the time the Join tool will try its best to join the two surfaces based on the selected edges from those surfaces. But there are times when joining doesnâ← t happen. Generally this occurs when the selected control-points are not completely describing the edge/row that you want to join. Select more control-points until the edge is completely highlighted.

Note that the edges do not have to be outside edges. You can join inside edges, although this is not typically done.

Skinning Section
This document was originally from a tutorial section. A copy of it was made and put in this section because the main index of the user manual for Blender did not seem to cover skinning for surfaces/meshes/curves. This document is out of date and needs some corrections but should work for surface skinning.

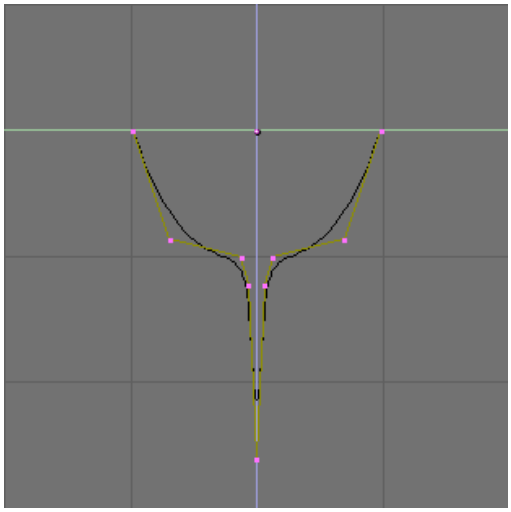
Skinning

Skinning is the fine art of defining a surface using two or more profiles. In Blender you do so by preparing as many curves of the the desired shape and then converting them to a single NURBS surface.



A Surface curve for skinning.

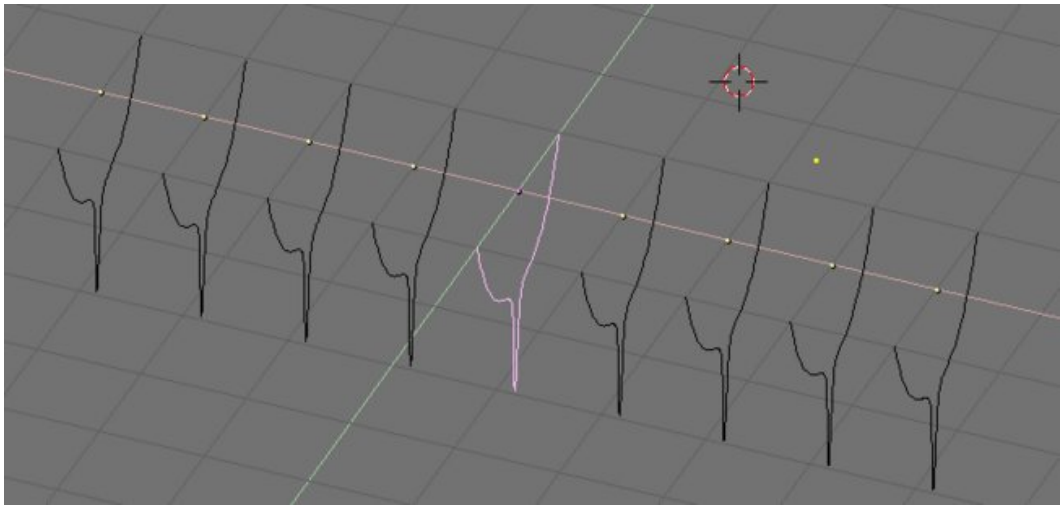
As an example we will create a sailboat. The first thing to do, in side view (3 NumPad), is to add a Surface Curve. Be sure to add a *Surface* curve and not a curve of BÃ©zier or NURBS flavour, or the trick wonâ t work (*A Surface curve for skinning*).



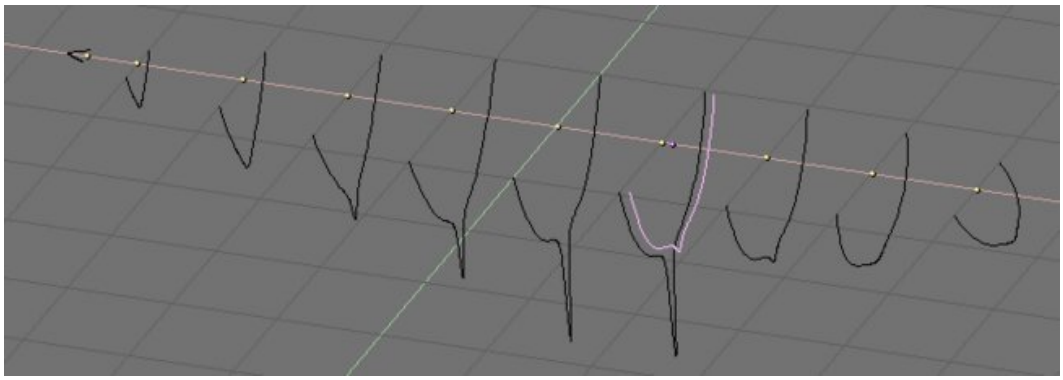
Profile of the ship.

Give the curve the shape of the middle cross section of the boat, by adding vertices as needed with the Split button and, possibly, by setting the NURBS to Endpoint both on â U â and â V â (*Profile of the ship*) as needed.

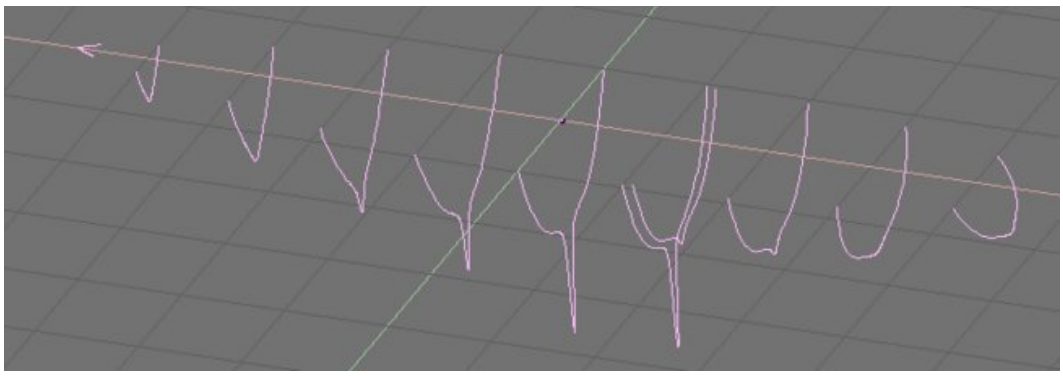
Now duplicate (â § ShiftD) the curve as many times as necessary, to the left and to the right (*Multiple profiles along shipâ s axis*). Adjust the curves to match the various sections of the ship at different points along its length. To this end, blueprints help a lot. You can load a blueprint on the background (as we did for the logo design in this chapter) to prepare all the cross section profiles (*Multiple profiles of the correct shapes*).



Multiple profiles along shipâ← s axis.
Note that the surface weâ← ll produce will be smoothed from one profile to the next. To create abrupt changes you would need to place profiles quite close to each other, as is the case for the profile selected in (*Multiple profiles of the correct shapes*).



Multiple profiles of the correct shapes.
Now select all curves (with A or B), and join them by pressing CtrlJ and by answering â← Yesâ← to the question â← `Join selected NURBS?`â← . The profiles are all highlighted in (*Joined profiles*).

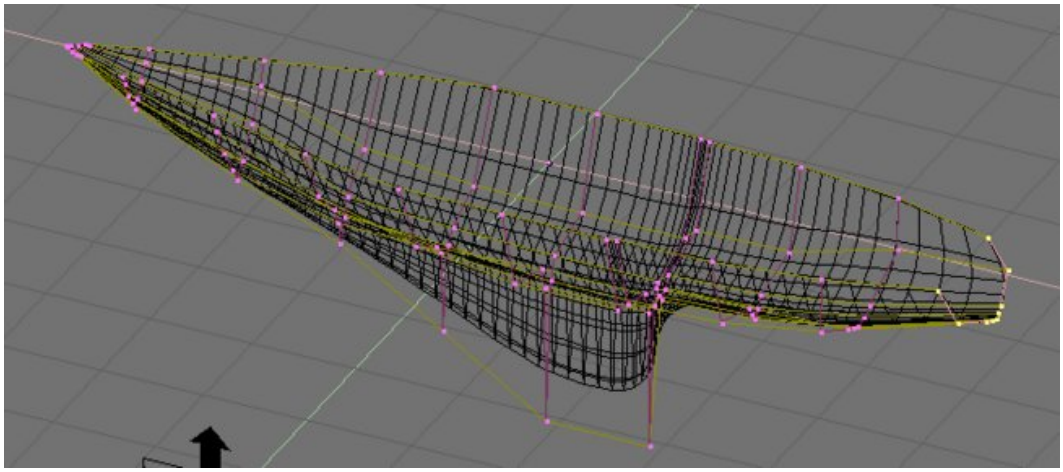


Joined profiles.
Now switch to Edit mode (â Tab) and select all control points with A; then press F. The profiles should be â← skinnedâ← and converted to a surface (*Skinned surface in Edit mode*).

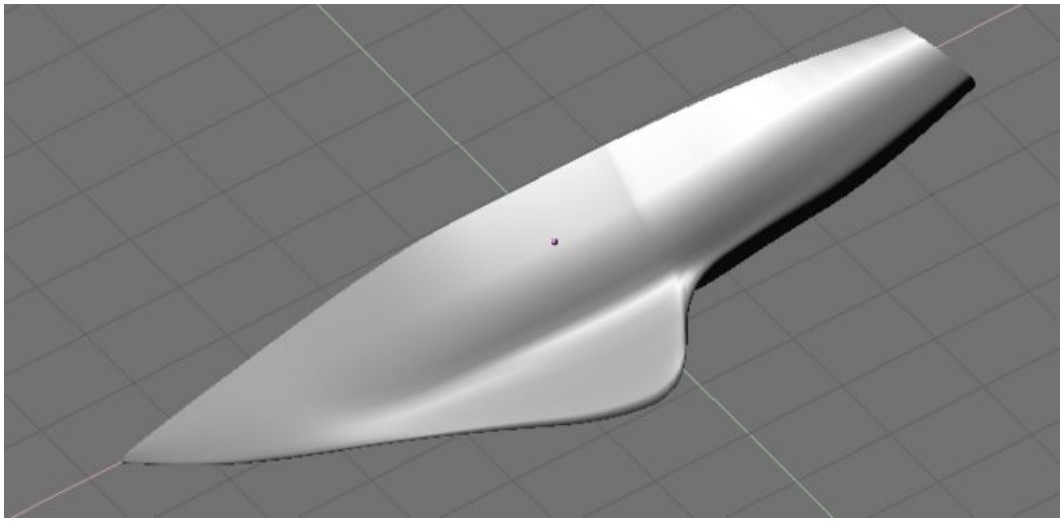
Note

As should be evident from the first and last profiles in this example, the cross-sections need not be defined on a family of mutually parallel planes.

WARNING
In some cases (i.e. when joining only two curves) the resulting surface may NOT be displayed after you press F. So you need to click on *Endpoint U/V* or *Uniform U/V* buttons on *Curve Tools* panel to correct this issue.



Skinned surface in Edit mode.
Tweak the surface, if necessary, by moving the control points. (*Final hull*) shows a shaded view. You will very probably need to increase ResolU and ResolV to obtain a better shape.



Final hull.

Profile setup
The only limitation to this otherwise very powerful technique is that all profiles must exhibit the same number of control points. This is why it is a good idea to model the most complex cross section first and then duplicate it, moving control points as needed, without adding or removing them, as weâ← ve shown in this example.

Text

Mode: Edit mode (Text)

Panel: Editing Context â Font

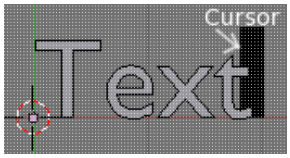
Hotkey: F9

Menu: Add â Text

Description

Text is considered a special Curve type that is completely separate from any of the other Curve types. Not only does the **Font system** have its own *built-in* font but it can use external fonts too, including *PostScript Type 1*, *OpenType* and *TrueType* fonts.

Options



Created text.
Creating a Text object is simple, use Add â Text. Once the text is created you are placed in Edit mode with the word â← `Text`â← inserted as a default placeholder, see (*Created text*). The *Black* block is the cursor.
created by W.S.F Corp (williamsouthall-ford)

Examples



Text Examples.
(*Text Examples*) shows some examples of various fonts in action including the â← blueâ← font that has been applied to a curve path.

Hints

A maximum of **50000** characters is allowed per text object, however, be forewarned that the more characters a single text object has, the slower the object will respond interactively.

Editing Text

Mode: Edit mode (Text)

Hotkey: see below

Description

Editing Text is similar to using a standard text editor but is not as full featured and is slightly different. To edit text, make sure the mouse cursor is placed over the graphical screen in which text is displayed.

Options

Exit Edit mode
: â Tab doesnâ← t insert a tab character in the text, but rather enters and exits Edit mode, as with other object types.

Copy
: To copy text to the buffer use CtrlC.

Cut and Copy
: To cut and copy text to the buffer use CtrlX.

Paste
: To paste text from the buffer use CtrlV.

Delete all text
: To completely erase or delete all text use Ctrlâ Backspace.

Home/End
: â Home and â ¥ End move the cursor to the beginning and end of a line respectively.

Next/Previous word
: To move the cursor on a wordâ← s boundary use Ctrlâ or Ctrlâ .

The text buffer does not communicate with the desktop. It only works from within Blender. To insert text from outside Blender see Inserting text.

Selecting text consists of holding down the â § Shift while using the **Arrow** keys (â /â /â /â) or Page up / Page down keys. The selection is remembered even in Object mode.

Inserting Text

You can insert text in three different ways: from the internal text buffer (Editing Text), with the â← Loremâ← button (used to insert at current position a piece of â← *Lorem Ipsum*â← which is bogus Latin text whose function is to see how text will look in a layout (also called "printer's Greek") â← yes, the same thing used by web developers!), or from a text file.

To load text from a text file click the Insert Text button on the Font panel. This will bring up a File Browser window for navigating to a valid UTF-8 file. As usual, be careful that the file doesnâ← t have too many characters as interactive response will slow down.

Changing Fonts

Mode: Edit mode (Text)

Panel: Editing Context â Font

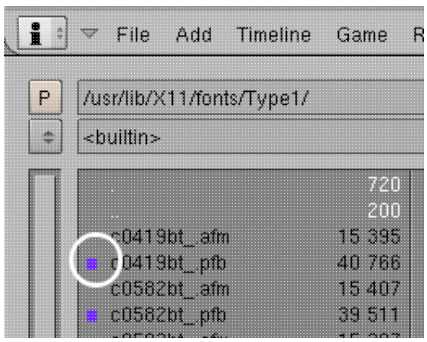
Hotkey: F9

Description



Font drop-down list and button.
Blender comes with a *built-in* font by default and is displayed in the drop-down list next to the Load button on the Font panel. The *built-in* font is always present and shows in this list as â← `<builtin>`â← .

Options



Loading a Type 1 font file.
To use a different Font you need to load it first by clicking the Load button in the Font panel and navigating to a *valid* font. The File Browser window will highlight any valid fonts by placing a small purplish rectangle next to each valid entry as shown in (*Loading a Type 1 font file*). The white circle highlights an example of a valid font.

Unix note
Fonts are typically located under `/usr/lib/fonts`, or some variant like `/usr/lib/X11/fonts`, but not always. They may be in other locations as well, such as `/usr/share/local` or `/usr/local/share`, and possibly related sub-trees.

If you select a font that Blender canâ ← t understand, you will get the error â ← `Not a valid fontâ ←` .

A separate font is required for each style. For example, you will need to load an *Italics* font in order to make characters or words italic. Once the font is loaded you can apply that font â ← Styleâ ← to the selected characters or the whole object. In all, you would need to load a minimum of four different types of fonts to represent each style (**Normal**, **Italics**, **Bold**, **Bold-Italics**).

Typography

Mode: Edit mode (Text)

Panel: Editing Context â Font

Hotkey: F9

Description

Blender has a number of typographic controls for changing the style and layout of text.

Options

Bold, Italics and Underline

Italics
: Toggled with CtrlI, font set with the I button.

Bold
: Toggled with CtrlB, font set with the B button.

Underline
: Toggled with CtrlU or by using the U button.

Blenderâ ← s B and I buttons donâ ← t work the same way as other applications. They serve as placeholders for you to load up certain fonts manually, which get applied when you use CtrlB or CtrlI when editing text.

To apply the Bold/Italics/Underline attribute to a set of characters you either turn on Bold/Italics/Underline prior to typing characters, or highlight first and then toggle Bold/Italics/Underline with a hotkey. Bold/Italics are applied based on a loaded font. For example, some characters may have one font representing normal characters and the `<builtin>` font representing Bold; see (*Bold text*).

Basically each font style is represented by a loaded font. One font may represents Bold while another font represents Italics (i.e. one font per style).

Alignment

Left
: Aligns text to left of frames when using them, else uses the centre point of the Text object as the starting point of the text (which grows to the right).

Center
: Centres text in the frames when using them, else uses the centre point of the Text object as the mid-point of the text (which grows equally to the left and right).

Right
: Aligns text to right of frames when using them, else uses the centre point of the Text object as the ending point of the text (which grows to the left).

Justify
: Only flushes a line when it is **terminated** by a wordwrap (**not** by â µ Enter), it uses *whitespace* instead of *character spacing* (kerning) to fill lines.

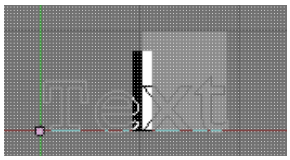
Flush
: **Always** flushes the line, even when itâ← s still being entered, it uses character spacing (kerning) to fill lines.

Both Justify and Flush only work within frames.

Word spacing
: A factor by which whitespace is scaled in width.

Spacing (kerning)
: Manual *kerning*, between all pairs of characters. You can also control it, **independently for each pair**, by pressing Altâ or Altâ to decrease/increase kerning by steps of **0.1**.

Examples



Bold text.
In (*Bold text*), one font is used for â← `Te`â← and a different font for â← `xt`â← .

Text Frames

Mode: Object mode / Edit mode (Text)

Panel: Editing Context â Font

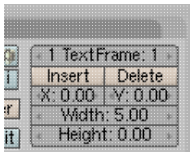
Hotkey: F9

Description

Text â← Framesâ← allow you to distribute the text amongst rectangular areas within a single text object. An arbitrary number of freely positionable and re-sizable text frames are allowed per text object.

Text flows continuously from the lowest-numbered Frame to the highest-numbered Frame with text inside each frame word-wrapped. Text flows between Frames when a lower-numbered frame canâ← t fit anymore text. If the last Frame is reached text overflows out of it.

Options



Text Frame.
Frames are controlled from the upper right corner of the Font panel; see (*Text Frame*).

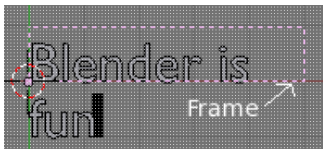
Frame size

By default the first Frame for a new text object, and any additional Frames, has a size of **Zero** for both Width and Height, which means the Frame is initially not visible.

Frames with a width of **0.0** are ignored completely during text flow (no wordwrap happens) and Frames with a height of **0.0** flow forever (no flowing to the next text frame).

In order for the frame to become visible the Frameâ← s Width must be greater than **0.0**.

Note
Technically the height is never actually **0.0** because the font itself always contributes height.



Frame width.
(*Frame width*) is a text object with a width of **5.0**. And because the frame width is greater than **0.0** it is now visible and is drawn in the active theme colour as a dashed rectangle. The text has overflowed because the text has reached the end of the last frame, the default frame.

Adding/Deleting a Frame

To add a Frame click the Insert button on the Font panel. A new frame is added with a default Width and Height of **0.0** which means it is not visible, nor will text flow from it into another frame. Be sure to set an offset for the new frame in the X and Y fields. Just an X setting will create a new column.

To delete the current Frame click the Delete button on the Font panel. Any text in higher frames will be re-flowed downward into lower frames.

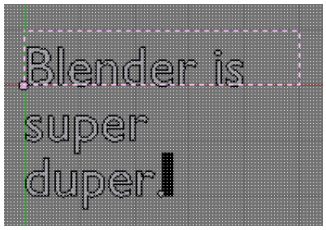
Examples

Text Frames are very similar to the concept of *frames* from a desktop publishing application, like Scribus. You use frames to control the placement and flow of text.

Text Flow

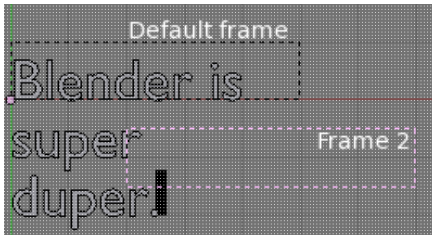


Text 1.
With two or more frames you can organize text to a finer degree. For example, create a text object and enter â← `Blender is super duper`â← ; see (*Text 1*). This text object has a frame, it just isnâ← t visible because its Width is **0.0**.



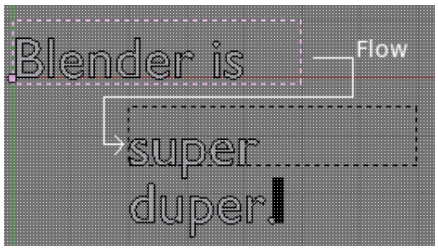
Text 2.
Set the width to **5.0**. The frame is now visible and text is wrapping according to the new width, as shown in (*Text 2*). Notice that the text has overflowed out of the frame. This is because the text has reached the end of the last frame which just happens to be the default/initial frame.

When we add another frame and set its width and height the text will flow into the new frame.



Text 3.
Clicking on Insert will add a new frame, labelled â← `Frame 2`â← in (*Text 3*); with the same attributes as the previous frame (labelled â← `Default frame`â← in (*Text 3*)).

Notice that the text has not yet flowed into this new frame. That is because the previous, or lower numbered, frame has a Height of **0.0**. Remember the height field may be **0.0** but the font itself contributes height. The fontâ← s height does not count. This means the height field value is an addition to the fontâ← s height.

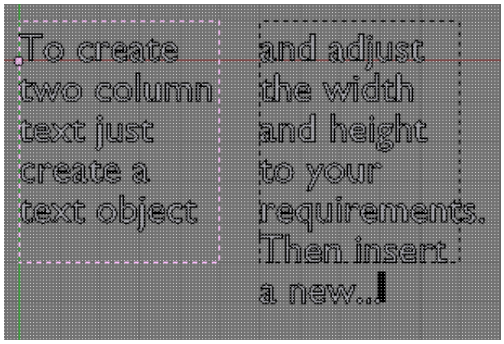


Text 4.

To get text to flow *into* â← Frame 2â← we need to change the height of the default/initial frame. In (*Text 4*) the height of the initial frame â← in pink â← has been increased to **0.1**.

Now the text flows from the initial frame into â← Frame 2â← . Notice that the text overflows out of â← Frame 2â← . Again this is because the text has reached the end of the last frame.

Multiple columns



Text 5.
To create two columns of text just create a text object and adjust the initial frameâ← s Width and Height to your requirements, then insert a new frame. The new frame will have the same size as the initial frame. Set the X position to something greater or less than the width of the initial frame; see (*Text 5*).

Multiple Materials

Mode: Object mode / Edit mode (Text)

Panel: Editing Context â Link and Materials

Hotkey: F9

Description

Each character can have a different Material index in order to have different materials on different characters.

Options

You can assign indices either as you type, or after by selecting blocks of text and clicking on the Assign button in the Link and Materials panel.

Examples



Red Green Blue.

For example to create (*Red Green Blue*) you would need to create three separate Materials and three separate Material indices. Each word would be assigned a Material index by selecting the characters for each word and clicking the Assign button. (*Red Green Blue*) is still one single Text object.

Curve and Surface attributes

Mode: Object mode / Edit mode (Text)

Panel: Editing Context â Curve and Surface

Hotkey: F9

Description

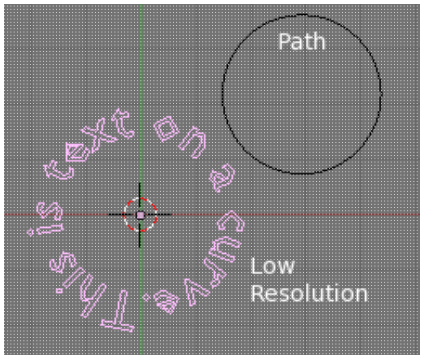
Text is very similar to 2D curves in that they have curve like properties. For example, you can change the Resolution from the Curve and Surface panel to have smooth or coarse text. Once the text is created you can extrude, bevel or even change the thickness.

Options

Because a Text object is similar to a curve it can be converted into a curve using AltC. Once this is performed the Text becomes a curve and can be manipulated just like a curve. This allows complete control over the shape of the characters beyond what a Text object provides.

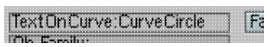
The transform from Text to Curve is not reversible; consider saving prior to converting. Also, you can continue to convert the Curve into a Mesh for even further control.

Examples



Rough text.
In (*Rough text*) the Resolution has been set to the lowest setting to produce very blocky text. Almost as if the text was broken out of a rock mine. In addition, the text has been applied to a 2D Bezier circle curve.

The path the text has been applied to is labelled â← `Path`â← . To specify a *Curve* or *Path* enter the name of a 2D curve in the TextOnCurve field in the Font panel as shown in (*TextOnCurve*). In this example the â← Pathâ← â← s name is â← `CurveCircle`â← .



TextOnCurve.

See Also

- Modelling: Curves
- Modelling: Surfaces.

Special Characters

Mode: Edit mode (Text)

Menu: Text â Special Characters

Description

There are a few special characters that are available using the Alt key or the Text menu on the 3D window header. These â← Keyâ← combinations are only available while in Edit mode.

Options

A summary of these characters follows; just remember you can access these characters from the Char panel as well:

AltC:	Copyright (Â©)	AltR:	Registered trademark (Â®)
AltG:	Degrees (Â°)	AltX:	Multiply symbol (Ã)
AltS:	German â← ssâ← (Ã)	AltF:	Currency sign (Â¤)
AltL:	British Pound (Â£)	AltY:	Japanese Yen (Â¥)
Alt1:	Superscript 1 (Â¹)	Alt2:	Superscript 2 (Â²)
Alt3:	Superscript 3 (Â³)		
Alt?:	Spanish question mark (Â¿)	Alt!:	Spanish exclamation mark (Â¡)
Alt<:	Left double quotation mark (Â«)	Alt>:	Right double quotation mark (Â»)

All the characters on your keyboard should work, including stressed vowels and so on. If you need special characters (such as accented letters, which are not there on a US keyboard) you can produce many of them using a combination of two other characters. To do so, type the main letter, press Altâ Backspace, and then press the desired â← modifierâ← to produce the special character. Some examples are given below :

A, Altâ Backspace, ~:	Ã£	A, Altâ Backspace, ':	Ã¡	A, Altâ Backspace, `:	Ã
A, Altâ Backspace, O:	Ã¥	E, Altâ Backspace, ":	Ã«	O, Altâ Backspace, /:	Ã¸

Unicode Characters

Mode: Edit mode (Text)

Panel: Editing Context â Char

Hotkey: F9

Description

The font system understands both ASCII and UNICODE character sets with a panel dedicated to assisting in the selection of extended characters.

Since Blender does not support Unicode text input via the keyboard, not all characters are easily accessible from the keyboard. For those difficult characters the Char panel is provided. This panel simply exposes the entire Unicode character set, by default. The character set can be quite large so paging buttons are provided, U and D. You can also select a subset of Unicode with the Unicode Table button, and jump directly to a known character by typing its *decimal* code in the num field just to the right of this button.

When you find the character you are looking for, just click on it in the grid.

Technical Details

For optimum resource usage only characters that are being used consume memory rather than the entire character set.

Meta Objects

Mode: Object Mode or Edit Mode (Meta)

Hotkey: â § ShiftA

Menu: Add â Meta

Meta Objects are *implicit surfaces*, meaning that they are *not explicitly* defined by vertices (as meshes are) or control points (as surfaces are): they exist *procedurally*. Meta objects are literally mathematical formulas that are calculated on-the-fly by Blender.

A very distinct visual characteristic of Meta Objects is that they are fluid *Mercurial*, or *Clay-like* forms that have a â← roundedâ← shape. Furthermore, when two meta objects get close to one another, they begin to interact with one another. They â← blendâ← or â← mergeâ← , as water droplets do (which, by the way, makes them very handy for modeling streams of water when you don't want to do a fluid simulation). If they subsequently move away from one another, they restore their original shape.

There are five predefined Meta Object configurations:

Ball
: A point underlying structure.

Tube

A line segment underlying structure.

Plane
A planar underlying structure.

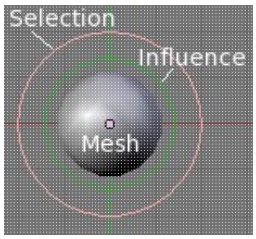
Elipsoid
A spherical underlying structure.

Cube
A volumetric cubic underlying structure.

Each of these is defined by its own underlying mathematical structure, and you can at any time switch between them using the MetaBall tools panel.

Typically Meta Objects are used for special effects or as a basis for modeling. For example, you could use a collection of Meta Objects to form the initial shape of your model and then convert it to another object type for further modeling. Meta objects are also very efficient for ray-tracing.

Options



MetaBall example.
Each Meta Object always appears with two rings or circles; see (*MetaBall example*).

Pink ring â← Selection

The outer ring (labeled â← `Selection`â← and coloured pink) is for *selecting* and exists because there are two types of elements that can be selected within a Meta Object. You can select the Meta Object itself, by clicking the ring, or select the Mesh by clicking the mesh. Without the selection ring you wouldnâ← t be able to select just the Meta Object.

S with the outer ring selected scales the Meta Object

Green Ring â← Stiffness

The inner ring (labeled â← `Influence`â← and coloured green) defines the Meta Objectâ← s stiffness value â← how much influence it has on other Meta Objects. See below for more details on this setting.

Technical Details

A more formal definition of a meta object can be given as a *directing structure* which can be seen as the source of a static field. The field can be either positive or negative and hence the field generated by neighbouring directing structures can attract or repel.

The implicit surface is defined as the surface where the 3D field generated by all the directing structures assume a given value. For example a Meta Ball, whose directing structure is a point, generates an isotropic (i.e. identical in all directions) field around it and the surfaces at constant field value are spheres centered at the directing point.

Meta Objects are nothing more than mathematical formulae that perform logical operations on one another (AND, OR), and that can be added and subtracted from each other. This method is also called **Constructive Solid Geometry** (CSG). Because of its mathematical nature, CSG uses little memory, but requires lots of processing power to compute.

Threshold (Influence)

Mode: Object Mode or Edit Mode (Meta)

Panel: Editing Context â MetaBall

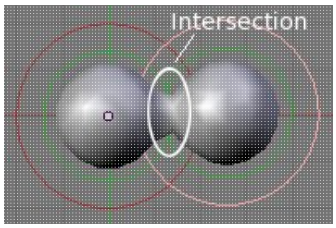
Description

Threshold defines how much a MetaObjectâ← s surface â← Influencesâ← other MetaObjects. It controls the *field level* at which the surface is computed. The setting is global to a Group of MetaObjects. As the Threshold increases, the influence that each MetaObject has on one another increases.

Options

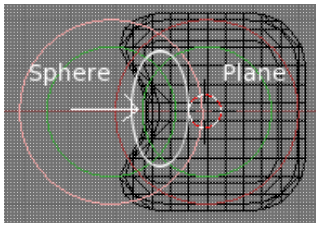
There are two types of influence: **positive** or **negative**. The type can be toggled on the MetaBall tools panel using the Negative button. You could think of **positive** as attraction and **negative** as repulsion of meshes. A negative MetaObject will push away or repel the meshes of positive MetaObjects.

â← Positiveâ← & â← Negativeâ← Influence



Positive.
A *Positive* influence is defined as an attraction meaning the meshes will stretch towards each other as the *rings of influence* intersect. (*Positive*) shows two MetaBallsâ← *ring of influence* intersecting with a *positive* influence.

Notice how the meshes have pulled towards one another. The area circled in white shows the green *influence* rings intersecting.



Negative.
The opposite effect of a *positive* influence would be a *negative* influence: the objects repel each other. (*Negative*) shows a MetaBall and MetaPlane where the MetaBall is negative and the MetaPlane is positive. Notice how the Negative MetaObject is not visible: only the surrounding circles appear. This is how Blender indicates that the object is Negative.

The white arrow indicates how the sphere is repelling or pushing away the planeâ← s mesh. This causes the planeâ← s mesh to â← cave inâ← or collapse inward. If you move the plane away from the sphere, the planeâ← s mesh will restore itself.

Very interesting screen-savers have been created using animations of positive and negative meta objects together in the same space!

Wiresize

Mode: Object Mode or Edit Mode (Meta)

Panel: Editing Context â MetaBall

Description

The Wiresize controls the resolution of the resultant mesh as generated by the MetaObject.

Options

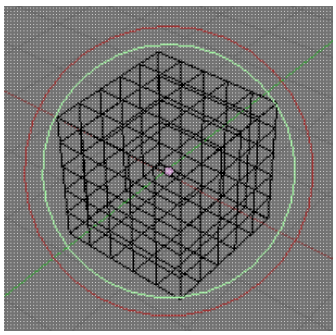
Wiresize
: The 3D View resolution of the generated mesh. The range is from **0.05** (finest) to **1.0** (coarsest).

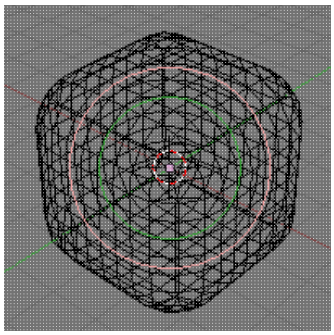
Rendersize
: The rendered resolution of the generated mesh. The range is from **0.05** (finest) to **1.0** (coarsest).

Examples

One way to see the underlying mathematical structure is to lower the Wiresize, increase the Threshold and set the Stiffness a fraction above the Threshold.

(*Underlying structure*) is a (*MetaCube*) with the above mentioned configuration applied as follows: Wiresize of **0.410**, Threshold of **5.0** and Stiffness a fraction above at **5.01**.


Underlying structure.


MetaCube.

You can clearly see the underlying Cube structure that gives the MetaCube its shape.

Stiffness

Mode: Edit Mode (Meta)

Panel: Editing Context â MetaBall tools

Hotkey: S

Description

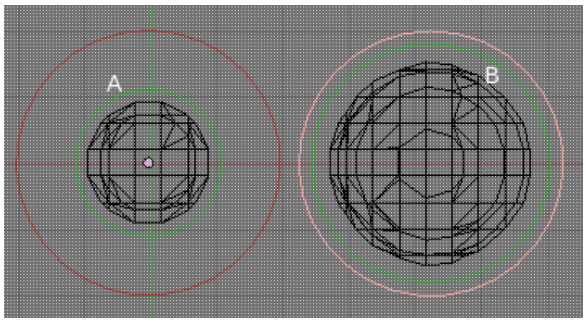
Together with Threshold, Stiffness controls the influencing range.

When a Meta Object comes within â← rangeâ← of another Meta Object the two Meta Objects will begin to interact with each other. They donâ← t necessarily need to intersect, and depending on the Threshold and Stiffness settings they most likely wonâ← t need to. Stiffness directly controls the *green ring* surrounding a MetaObject. The field is found on the MetaBall tools panel.

Options

The range is from **0.0** to **10.0**. But to be visible the Stiffness must be slightly larger than the Threshold value. You can visually adjust the Stiffness ring by using the RMB to select it and activating Scale mode with the S.

Examples



Stiffness.

In (*Stiffness*), the MetaBall labelled â← Aâ← , has a smaller Stiffness value than the MetaBall, labelled â← Bâ← . As you can see the *green ring* radius is different between them.

Grouping

Mode: Object Mode or Edit Mode (Meta)

Panel: Editing Context â Link and Materials

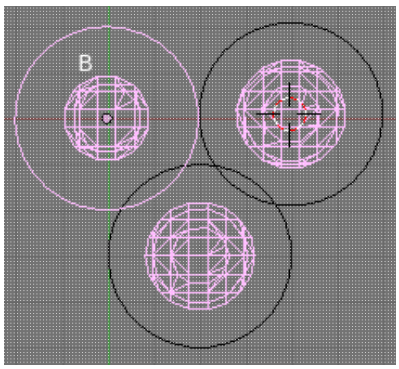
Hotkey: F9

Description

MetaObjects are grouped by the *Family* part of an Object name (this is the â← OB:â← field in most panels, **not** the â← MB:â← field).

The Object name is broken into two parts: the left part before, the period, and the right part, after the period. The left part is the â← family nameâ← , while the right part is simply a sequential number whose only purpose is to make the object-names unique. For example, the *Family* part of â← `MetaPlane.001`â← is â← MetaPlaneâ← . Each meta object in the same â← familyâ← is associated with one another as discussed below.

Options



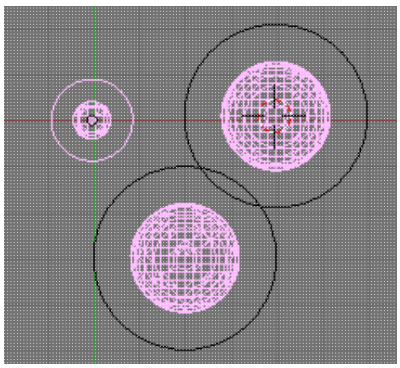
MetaBall Base.
Groups of MetaObjects are controlled by a *base* MetaObject which is identified by an Object name **without** a â← numberâ← part. For example, if we have five MetaObjects called â← `MetaThing`â← , â← `MetaThing.001`â← , â← `MetaThing.002`â← , â← `MetaThing.003`â← and â← `MetaThing.004`â← , the *base* MetaObject would be â← `MetaThing`â← .

The *base* MetaObject determines the basis, the resolution, *and* the transformations. It also has the material and texture area. The *base* MetaObject is effectively the parent of (or perhaps a better word to use is "the owner of") the other MetaObjects in the group.

A group can only have one Material and Texture area. This normalises the coordinates of the vertices. Normally the texture area is identical to the *bounding box* of all vertices. The user can force a texture area with the T command in Object mode.

Examples

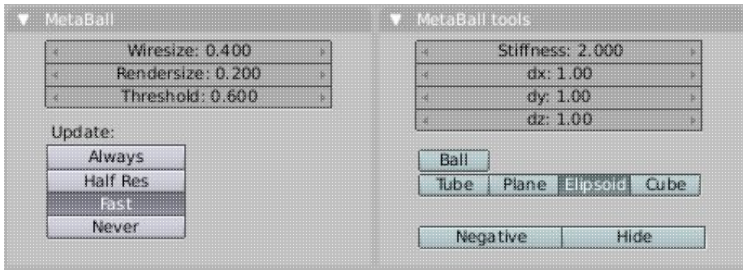
(*MetaBall Base*) shows the *base* MetaObject labelled â← Bâ← . The other two MetaObjects are *children*. Childrenâ← s selection rings are always black, while the Groupâ← s mesh is pink. Because the MetaObjects are grouped, they form a unified mesh which can always be selected by selecting the mesh of any MetaObject in the group. For example, in the example (*MetaBall Base*), only the lower sphere (the parent) has been selected, and you see that both the parent's mesh *and* all of the children's meshes are now highlighted.



Scaling the â← baseâ← .
The *base* MetaObject controls the **polygonalization** (mesh structure) for the group, and as such, also controls the polygonalization for the children (*non-base*) MetaObjects. If we transform the *base* MetaObject, the childrenâ← s polygonalization changes. However, if we transform the children, the polygonalization remains unchanged.

Hints

This discussion of â← polygonizationâ← *doesnâ← t* mean that the various meshes donâ← t deform towards or away from each other (Meta objects always influence one another in the usual way, whether or not they are members of the same family). Rather, it means that the underlying mesh structure changes only when the *base* object transforms. For example, if you Scale the *base*, the childrenâ← s mesh structure changes. In (*Scaling the â← baseâ←*), the *base* has been scaled down, which has the effect of scaling the mesh structure of each of the children. As you can see, the childrenâ← s mesh resolution has increased, while the *base* decreased. *The children did not change size!*



Metaballs use a different modelling workflow from more traditional techniques. You guide their shape rather then editing their vertices. Use the Editing context to guide their shape.

The MetaBall panel controls how they look in your 3D view, and how smooth they are when rendered. Choose a smaller Rendersize to make them smoother when rendering. For real-time updates while you are working with them inside Blender, choose a larger Wiresize. The Threshold sets how close they can get to each other before merging. The Update: selections choose your display update frequency. Use Fast to save CPU time and increase Blenderâ← s responsiveness to you. Use Never if you really want to get confused between your display and what is rendered.

The MetaBall tools panel controls their shape. Select Ball for a sphere, Tube, Plane, etc. Each choice is like the mesh equivalent with lots of subsurfacing.

Main Section for Scripts

There are two main sets of documentation of the Blender scripts:

- Extensions:Py/Scripts/Manual: Contains Tutorials on using Python. Links to script information, Blender Python API docs.
- Extensions:Py/Scripts/Catalog: A large listing of scripts, grouped by function, linking to individual script pages.

Search the Catalog to find and download the script you need.

- Some Modelling Script Samples:

Poly Reducer

Mode: Edit Mode (Mesh)

Menu: Mesh â Scripts â Poly Reducer

Description

This tool can be used instead of Blenders Decimate modifier as a way to remove polygons from a mesh while keeping the original shape as much as possible.

Reasons you may want to use a polygon reducer are:

- To make 3D Scanned data usable when rendering and editing.
- Generate Level Of Detail models (LODâ← s), for games or simulation models.
- To speed up render times.

Options

Poly Reduce is accessed from Edit Mode and will operate on the entire mesh.

On activation a popup will appear with the following options:

Poly Reduce
: Scale the meshâ← s poly count by this value.

Boundary Weight
: Weight boundary verts by this scale. Zero disables boundary weighting. A boundary vert is a vert that is not completely surrounded by faces. Some meshes have no boundary verts. For example, a cube has no boundary verts whereas a plane has only boundary verts.

Area Weight
: Collapse edges affecting lower area faces first. Zero disables area weighting.

Triangulate
: Convert quads to tris before reduction, for more choices of edges to collapse.
The advantage of triangulating is you have a larger set of edges to choose from when collapsing giving a higher quality result.

UV Coords
: Interpolate UV Coords (if existing).

Vert Colors
: Interpolate Vertex Colors (if existing).

Vert Weights
: Interpolate Vertex Weights (if existing).

Hints

Poly Reducer has some advantages and disadvantages compared to Blenderâ← s Decimate modifier, here are some pros and cons.

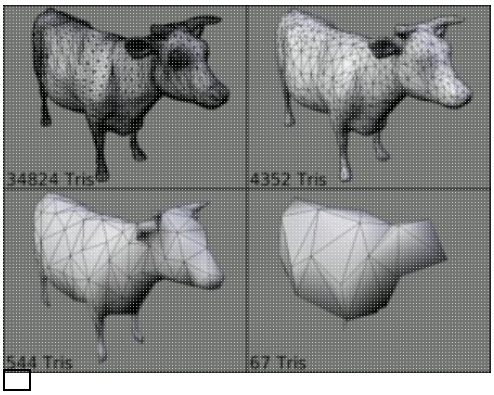
Pros

- Higher quality resulting mesh.
- Can operate on any mesh, will not throw errors if the mesh has odd face/edge/vert topology.
- Options to control where polygons are removed.
- Keeps materials assigned to faces.
- Maintains UV Texture coordinates, Vertex colors, and Vertex Group Weights (used for bone weight painting) - This makes it very useful for game/realtime models.

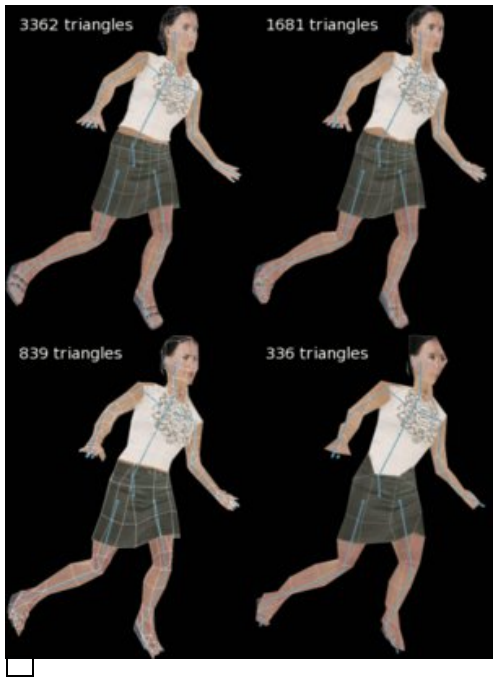
Cons

- Fairly Slow.
- Uses a lot of memory.

Examples



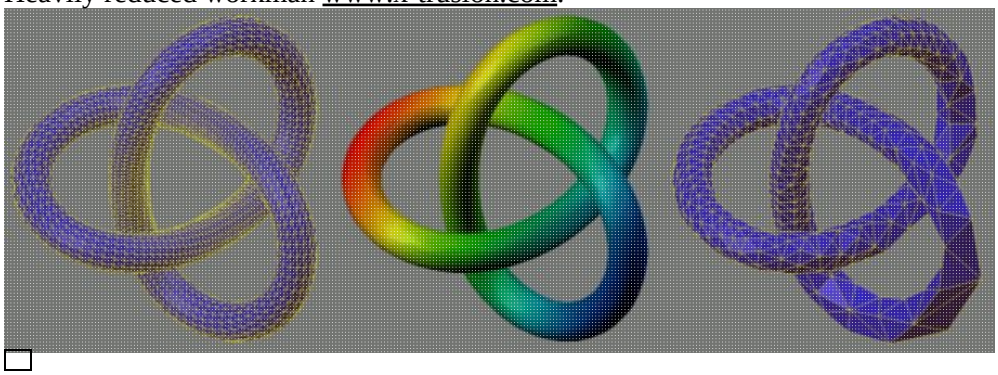
Famous cow.



Human with UV textures and bone weights from www.x-trusion.com.



Heavily reduced workman www.x-trusion.com.



Example of an 80% Reduction using a weight map for influencing the result. Original, Weight Map, Result of the reduction.

Auto Image Layout

Mode: All Modes (Mesh)

Menu: UV/Image Editor â UVs â Auto Image Layout

Description

This script makes a new image from the used areas of all the images mapped to the selected mesh objects.

Images are packed into 1 new image that is assigned to the original faces.

This is useful for game models where 1 image is faster than many, and saves the labour of manual texture layout in an image editor.

Options

This script is accessed from UV/Face mode and packs images from the active mesh.

On activation a popup will be appear with the following options.

image path
: No ext!
A new PNG image file will be created at this path. Use `//` as a prefix for the current blend file location. Otherwise you may specify the full path. Do not add in a file extension.

Pixel Size
: The size of the image, this value is used for width and height to make a square image.

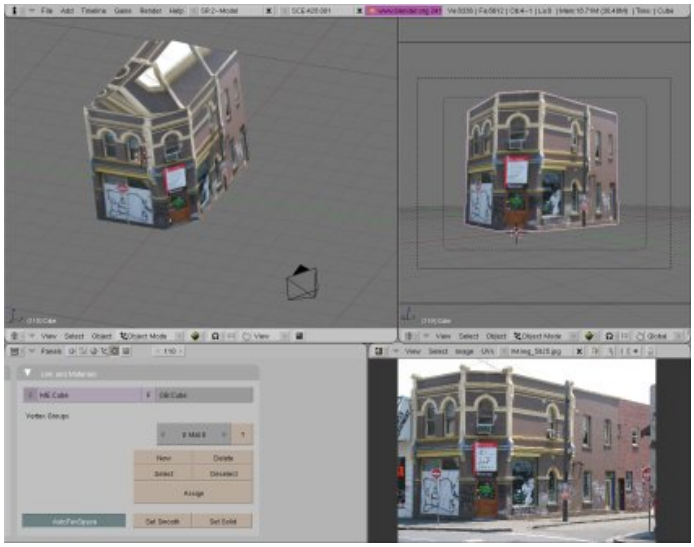
Pixel Margin
: When cropping the image to the bounds of the used areas add this pixel margin, this stops lower resolution textures (mipmaps) from bleeding the edge colour into the faces that use this texture.

Keep Image Aspect
: If this is turned off, the tiles will stretch to the bounds of the image, making the images look stretched in an image viewer. However it will give better results when viewed in 3d because there is more pixel information in the image.

Texture Source, All Sel Objects
: When enabled all selected objects will have their textures packed into the texture.

Examples

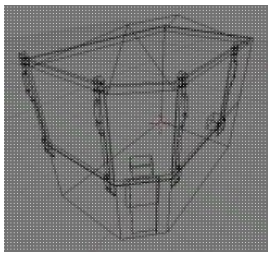
Here is a test case where I took 5 unedited photos, mapped them to a low poly mesh, and pack them into 1 texture.



Projection mapped uv mesh.



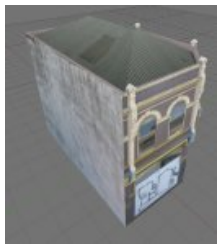
Model in details.



Finished Details with roof and side walls.



Texture view.



Back wall with generic texture.







Result of running the Auto Texture Layout Script.







All images used for this mesh.

Weight Painting

Bone Weight Copy

Mode: Object Mode (Mesh)

Menu: Object â Scripts â Bone Weight Copy

Description

This copies weights from one mesh to another based on vertex locations. It can also be used to update a mesh thatâ← s already weighted, by selecting the verts on the target mesh. Then using the Copy To Selected option.

Mirror Vertex Locations & Weight

Mode: Edit Mode (Mesh)

Menu: Mesh â Scripts â Mirror Vertex Locations & Weight

Description

This script is used to mirror vertex locations and weights. It is useful if you have a model that was made symmetrical but has verts that have moved from their mirrored locations slightly, causing Blenderâ← s X-Mirror options not to work.

Weights can be mirrored too, this is useful if you want to model 1 side of a mesh, copy the mesh and flip it. You can then use this script to mirror to the copy, even creating new flipped vertex groups, renaming group name `.left` to `.right` or `.L` to `.R`.

Vertex positions are mirrored by doing a locational lookup, finding matching verts on both sides of a mesh and moving to the left/right or mid location.

The vertex weights work differently, they are mirrored by location also, but they donâ← t mirror in pairs, rather it works by finding the closest vertex on the flip side and using its weight.

When a location mirror is finished, verts that have not been mirrored will remain selected. A good way to check both sides are mirrored is to select the mirrored parts, run this script with default options and then see if there are any selected verts.

For details on each option read the tooltips.

Weight Paint Gradient

Mode: Weight Paint (Mesh)

Menu: Paint â Weight Gradient

Description

Mix weight paint and face select mode so as to select the faces to gradient. Then Run Weight Gradient from the Paint menu, and click on the 2 locations to blend between. The existing weight under the mouse is used for to/from weights.

Vertex Colour Gradient

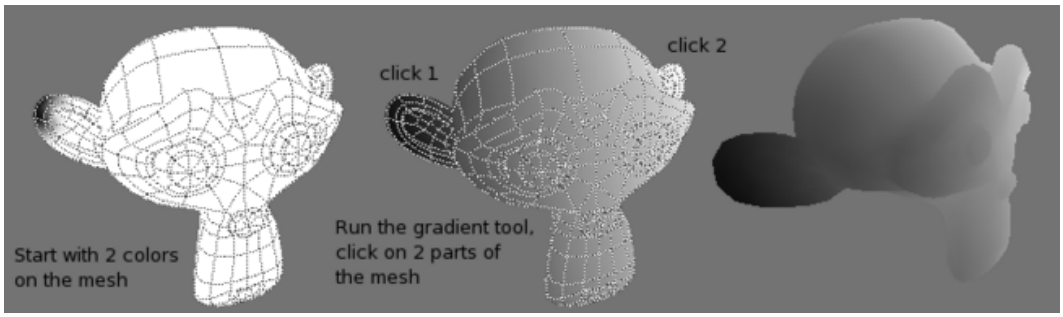
Mode: Vertex Paint (Mesh)

Menu: Paint â VCol Gradient

Description

See Weight Paint Gradient aboveâ←¦

Examples



Example of gradient usage.

Self Shadow

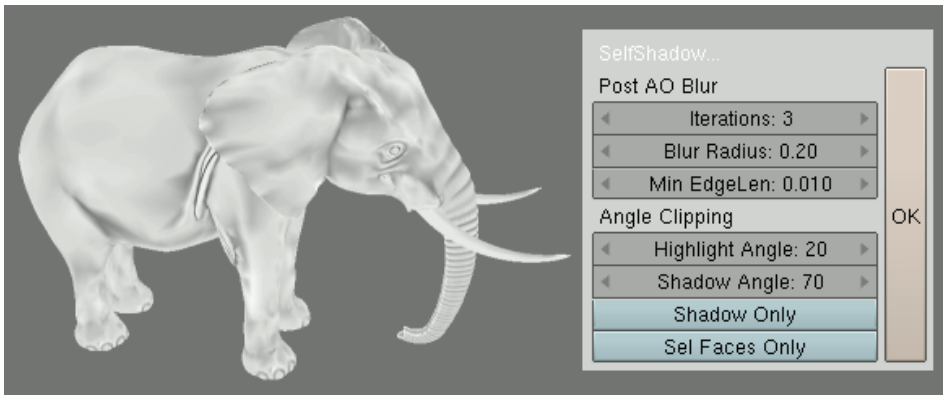
Mode: Vertex Paint (Mesh)

Menu: Paint â Self Shadow VCols (AO)

Description

Uses the mesh geometry to shade the mesh, similar to Ambient Occlusion.

Examples



Elephant Shaded.

Modifiers and deformation

Modifiers are automatic operations that affect an object in a non-destructive way. With modifiers, you can perform many effects automatically that would be otherwise tedious to do manually (such as subdivision surfaces) and without affecting the base topology of your object. Modifiers work by changing how an object is displayed and rendered, but not the actual object geometry. You can Apply a modifier if you wish to make its changes permanent.

There are three types of modifiers:

Deform modifiers
: They only change the shape of an object, and are available for meshes, and often texts, curves, surfaces and/or lattices.

Constructive modifiers
: They are only available for mesh objects, and typically change mesh topology in some way (e.g. the Subsurf modifier which performs Catmull-Clarke subdivision).

Virtual modifiers
: Currently, they only regard meshes. They have the specificities of having no settings on their own, and of *not being directly addable from the Modifiers panel* (they are automatically added when you activate the tool they represent). Their only role is to define the place in the modifiers stack used as base data by the tool they represent.

Interface

The interface that is used for modifiers and constraints is described here:

- Modifier Stack

Modifiers

Mode: Any mode

Panel: Modifiers (Editing context, F9)

Modifiers are added from the Modifiers tab in the Editing context (F9). The Modifiers tab appears when a Mesh, Curve, Surface, Text, or Lattice object is added or selected.

Deform modifiers

- **Armature** â← Use bones to deform and animate your object.
- **Cast** â← Shift the shape of a mesh, surface or lattice to a sphere, cylinder or cuboid.
- **Curve** â← Bend your object using a curve as guide.
- **Hooks** â← Add a hook to your vertice(s) (or control point(s)) to manipulate them from the outside.
- **Lattice** â← Use a Lattice object to deform your object.

- **MeshDeform** â← Allows you to deform your object by modifying the shape of another mesh, used as a â← Mesh Deform Cageâ← (like when using a lattice).
- **Shrinkwrap** â← Allows you to shrink/wrap your object to/around the surface of a target mesh object.
- **SimpleDeform** â← Applies some advanced deformations to your object.
- **Wave** â← Deform your object to form (animated) waves.

Constructive modifiers

- **Array** â← Create an array out of your basic mesh and similar (repeating) shapes.
- **Bevel** â← Create a bevel on a selected mesh object.
- **Booleans** â← Combine/subtract/intersect your mesh with another one.
- **Build** â← Assemble your mesh step by step when animating.
- **Decimate** â← Reduce the polygon count of your mesh.
- **Displace** â← Use textures or objects to displace your mesh.
- **EdgeSplit** â← Add sharp edges to your mesh.
- **Explode** â← Splits apart a mesh (to be used with particles).
- **Mask** â← Allows you to hide some parts of your mesh.
- **Mirror** â← Mirror an object about one of its own axis, so that the resultant mesh is symmetrical, and you only have to model/edit half, a fourth or a eighth of it.
- **Particle Instance** â← Make an object act similar to a particle but using the mesh shape instead.
- **Smooth** â← Smooth a mesh by flattening the angles between its faces.
- **SubSurf** â← Smooth the surface by creating interpolated geometry.
- **UVProject** â← Project UV coordinates on your mesh.

Virtual modifiers

- **Cloth** â← Simulates the properties of a piece of cloth. It is inserted in the modifier stack when you designate a mesh as Cloth.
- **Fluidsim** â← The object is part of a fluid simulationâ←¦ Modifier added when you designate a mesh as Fluid. Note that oddly, this modifier cannot be removed from the stack from the Modifiers panelâ←¦
- **Particle System** â← Represents a particle system in the stack, so it is inserted when you add a particle system to the object.
- **Softbody** â← The object is soft, elasticatedâ←¦ Modifier added when you designate a mesh as Softbody.

Modifier

A modifier is defined as the application of a â← processâ← or â← algorithmâ← upon objects. They can be applied interactively and non-destructively in just about any order the users chooses. This kind of functionality is often referred to as a â← modifier stackâ← and is found in several other 3D applications.

Modifiers are added from the Modifiers tab in the Editing context (F9). The Modifiers tab appears when a Mesh, Curve, Surface, Text, or Lattice object is added or selected.

Some â← toolsâ← , for example the â← Decimatorâ← , have been migrated from its previous location and changed into a modifier.

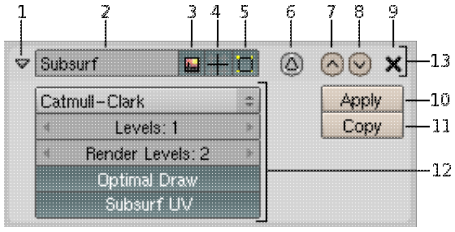
In a modifier stack the order in which modifiers are applied has an effect on the result. Fortunately modifiers can be rearranged easily by clicking the convenient up and down arrow icons. For example, (*Stack ordering*) shows SubSurf and Mirror modifiers that have switched places.

The left side of the yellow line has the Mirror modifier as the last item on the stack. In the lower half you can see two spheres. One is a â← mirrorâ← of the other. The original sphere is on the right and the â← mirrorâ← on the left. The Mirror modifierâ← s Merge Limit: has been set to a value of **0.312** to cause the vertices to weld together from a greater distance. The area marked by a white circle is a suggested area to concentrate on when the stack order is changed later.

The right side of the yellow line has the Subsurf modifier â← switched/rearrangedâ← to the bottom of the stack (i.e. it has switched places with the Mirror modifier). In this case the Subsurf modifier is being applied last. Now take a look at the white circle area. You can see that the results look different from the previous. This means that the stack order is very important in defining the end results.

Interface

Each modifier has been brought in from a different part of Blender, so each has its own unique settings and special considerations. However, each modifierâ← s interface has the same basic components, see (*Panel Layout (Subsurf as an example)*).



Panel Layout (Subsurf as an example).

- **1** â←  â← Collapses modifier to show only the header.
- **2** â← A name for this modifier â← Â the default being the name of the modifier itself. It is unique amongst other modifiers of the same type.
- **3** â←  â← Shows modifier effect in the rendering view.
- **4** â←  â← Shows modifier effect in the 3D view.
- **5** â←  â← Shows modifier effect in Edit mode. This button may not be available depending on the type of modifier.
- **6** â← , ,  â← Applies modifier to editing cage in Edit mode. The icon can be â← Disabledâ← , â← Deactivatedâ← and â← Activatedâ← , respectively. This icon materialize the Cage mode.
- **7** â←  â← Moves modifier up in the stack.
- **8** â←  â← Moves modifier down in the stack.
- **9** â←  â← Removes the modifier from the stack.
- **10** â← Apply â← Makes the modifier real.
- **11** â← Copy â← Creates a copy of the modifier at the base of the stack.
- **12** â← Sub panel for individual modifiers.
- **13** â← Header area for the main modifier controls.

Every modifier features a Collapse Arrow (**1**) and Name Box (**2**). The Collapse Arrow hides the modifierâ← s sub-panel (**12**) so that multiple modifiers can be displayed without the need for excessive scrolling in the buttons window. When collapsed, only the modifierâ← s header (**13**) is displayed. The Name Box can be used to give your modifiers titles, which can make recognizing their functions easier. This comes in handy in large scenes with complex modifier setups that feature multiple modifier types.

They are followed by three buttons which control the visibility of the effect in three separate contexts: Rendering (**3**), Object mode (**4**), and Edit mode (**5**). Toggling each button determines whether the modifierâ← s result displays in each mode.

The Cage mode (**6**) button is used to apply the modifier to the editing cage, which generates a more accurate display of the underlying geometry once a modifier has been applied. So itâ← s only available when the visibility in Edit mode is activated (**5**). This displays vert/edge/face positions in their â← modifiedâ← locations instead of their original locations. It should be noted, however, that transformation operations still act on the original locations of the cage vertices and not on the displayed locations. This button has three states: â← Disabledâ← , â← Activatedâ← and â← Deactivatedâ← . However, you can only control the tow last states by clicking on this control. When you activate it for a given modifier, all those above this one in the stack, and activated in Edit mode, are put in disabled Cage mode: you cannot have a modifier with the Cage mode deactivated, followed by one with the Cage mode activated!

The two Arrow buttons (**7** and **8**) control the order in which modifiers exist in the stack. Modifiers are evaluated top to bottom in the panel. The higher in the panel, the earlier it is evaluated. This can be very important depending on the application.

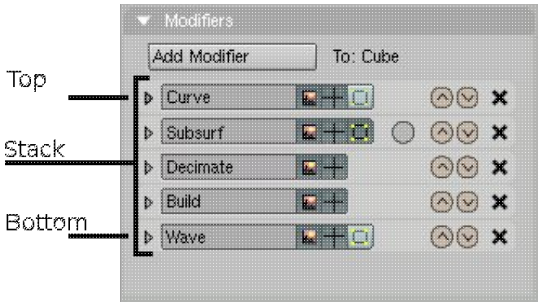
A great example of a situation in which positioning in the stack is important is the use of the Subsurf modifier in combination with the Mirror modifier as shown in (*Stack Ordering*) from the previous section. The Mirror modifier must appear before the SubSurf modifier in the list, otherwise the original half of the object gets subsurfaced and **then** mirrored, which may not be what you expected.

The Delete (**9**) button does exactly what one would expect, it removes the modifier from the stack entirely.

The Apply (**10**) and Copy (**11**) buttons have two very different functions despite their proximity to each other. Apply evaluates the modifier **as if it were the first modifier in the stack** and writes the results into the mesh, in effect â← bakingâ← the result of that modifier into the object. The Copy button creates a copy of the current modifier, including its settings, at the bottom of the modifier list.

Stack

To add a modifier you add it to the *stack*. Once added (always at the bottom of the stack), they can be rearranged to your liking.



Stack Example.
There used to be in previous Blender versions (prior 2.49) modifiers that had to be on top of the stack (i.e. be the first evaluated), but this restriction has now disappeared.

Some modifiers can only be applied to certain object types. This is indicated by the panel filtering the Add Modifier button on the Modifiers panel. Only modifiers that can be applied are shown in this listbox button.

For example, Mesh objects can have all available modifiers applied. But Lattice type objects can only have Armature, Cast, Curve, Lattice, MeshDeform, Shrinkwrap, SimpleDeform and Wave modifiers applied.

Armature Modifier

Mode: Object mode

Panel: Modifiers (Editing context, F9)

Description

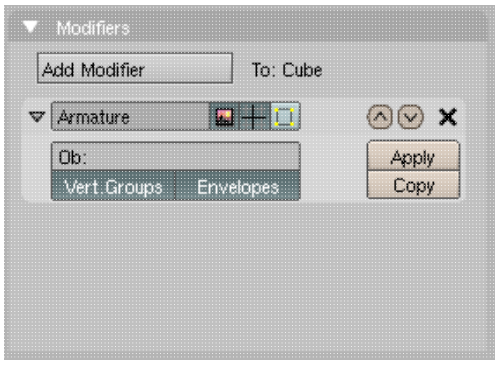
This modifier appears (as the others) as a sub-panel, in the Editing context which is accessed using F9 or clicking button in the Buttons window. This sub-panel is part of the Modifiers tab. For further information about the common panel components see the interface section on modifiers.

The Armature modifier is used for building skeletal systems for animating the poses of characters and anything else which needs to be posed. With the VGroup and MultiModifier options, you can use several armatures to animate a single (mesh) object.

By adding an armature system to an object, that object can be deformed accurately so that geometry doesnâ← t have to be animated by hand. The Armature modifier allows objects to be deformed by bones simply by specifying the name of the armature object, without having to use the (old) â← parent/childâ← system.

For more details on armatures usage, see this chapter.

Options



Modifiers panel with an Armature modifier activated.

Ob
: The name of the armature object used by this modifier.

VGroup
: The name of a vertex group of the object, of which weights will be used to determine the influence of this Armature modifierâ← s result when mixing it with the results from other Armature ones. Only meaningful when having at least two of these modifiers on the same object, with MultiModifier activated.

Inv
: Inverts the influence set by the vertex group defined in previous setting (i.e. reverts the weight values of this group).

Vert.Groups
: Enable/Disable vertex groups defining the deformation (i.e. bones of a given name only deform vertices belonging to groups of same name).

Envelopes
: Enable/Disable bone envelopes defining the deformation (i.e. bones deform vertices in their neighborhood).

Quaternion
: Use quaternions to get smoother and nicer rotation interpolations.

B-Bone Rest
: When using B-bones (â← curvedâ← bones), deform the object according to the B-bones â← curvesâ← , even when they are in rest position.

MultiModifier
: Use the same data as previous (Armature?) modifier as input. This allows you to use several armatures to deform a same object, all based on the â← non-deformedâ← data (i.e. this avoid the second Armature modifier to deform the result of the first oneâ←¦). The results of the Armature modifiers are then mixed together, using the weights of the VGroup vertex groups as â← mixing guidesâ← .

Cast Modifier

Mode: Any mode

Panel: Modifiers (Editing context, F9)

Description

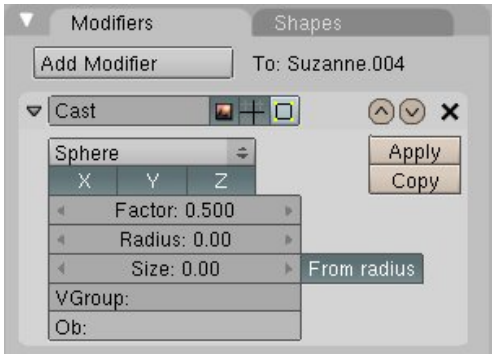
This modifier shifts the shape of a mesh, curve, surface or lattice to any of a few pre-defined shapes (sphere, cylinder, cuboid).

It is equivalent to the To Sphere button in the Editing contextÂ â Mesh Tools panel and the To Sphere command in the MeshÂ â Transform menu of the 3D views in Edit mode (â § ShiftCtrlS), as well as what other programs call â← Spherifyâ← or â← Spherizeâ← , but, as written above, it is not limited to casting to a sphere.

Hint: the Smooth modifier is a good companion to Cast, since the casted shape sometimes needs smoothing to look nicer or even to fix shading artifacts.

Important
For performance, this modifier works only with local coordinates. If the modified object looks wrong, you may need to apply the objectâ← s rotation (CtrlA), specially when casting to a cylinder.

Options



Cast modifier.

Type
: Menu to choose cast type (target shape): Sphere, Cylinder or Cuboid.

X, Y, Z
: Toggle buttons to enable/disable the modifier in the X, Y, Z axes directions (X and Y only for Cylinder cast type).

Factor
: The factor to control blending between original and casted vertex positions. Itâ← s a linear interpolation: **0.0** gives original coordinates (aka modifier has no effect), **1.0** casts to the target shape. Values below or above `[0.0, 1.0]` deform the mesh, sometimes in interesting ways.

Radius
: If nonzero, this radius defines a sphere of influence. Vertices outside it are not affected by the modifier.

Size
: Alternative size for the projected shape. If zero, it is defined by the initial shape and the control object, if any.

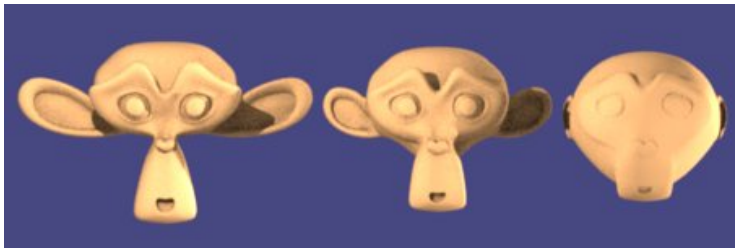
From radius
: If activated, calculate Size from Radius, for smoother results.

VGroup
: A vertex group name, to restrict the effect to the vertices in it only. This allows for selective, realtime casting, by painting vertex weights.

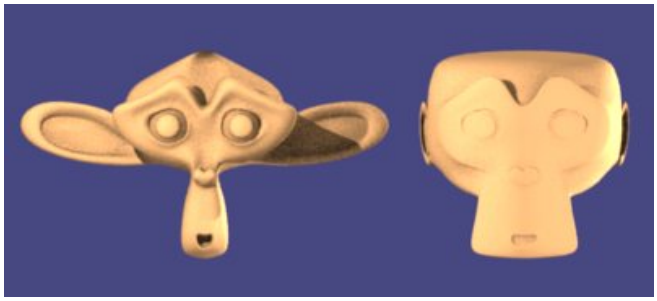
Ob
: The name of an object to control the effect. The location of this objectâ ← s center defines the center of the projection. Also, its size and rotation transform the projected vertices. Hint: animating (keyframing) this control object also animates the modified object.

Examples

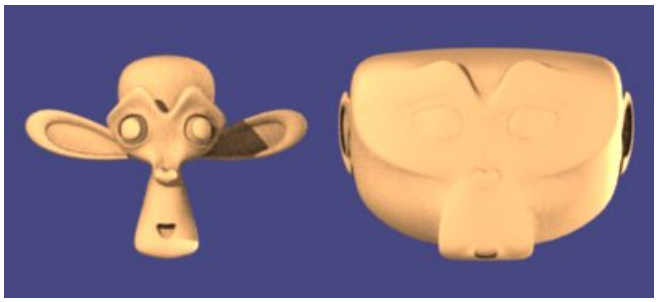
These models also have Subsurf modifiers applied after the Cast one:



Casting to sphere: **Left**: Factor: **-0.8**. **Center**: Factor: **0.0** (disabled). **Right**: Factor: **0.8**.

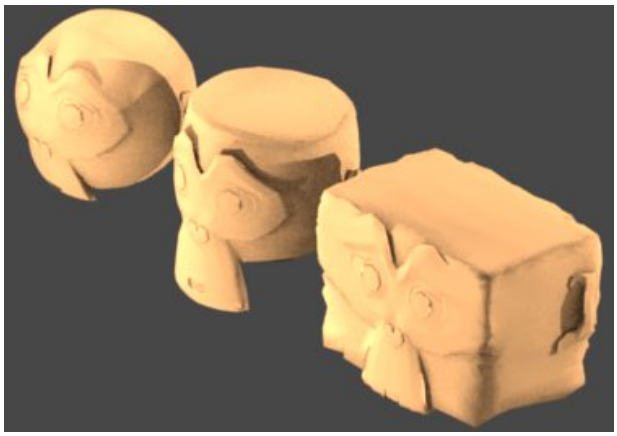


Casting to cylinder: **Left**: Factor: **-0.8**. **Right**: Factor: **0.8**.



Casting to cuboid: **Left**: Factor: **-0.8**. **Right**: Factor: **0.8**.

These have Subsurf (in Simple Subdiv mode) applied before the Cast modifier, to have more vertices to work with:



Casting Suzanne: **Left**: to Sphere; **Center**: to Cylinder. **Right**: to Cuboid.

Curve Modifier

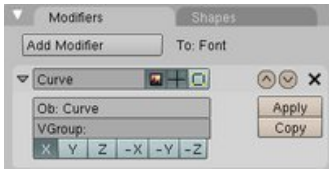
Mode: Object mode

Panel: Modifiers (Editing context, F9)

Description

The Curve modifier functions just like its predecessor with the added exception that there is no need for a parent/child relationship between the curve and the object being deformed, and that the effect can be applied to all object types in realtime.

Options



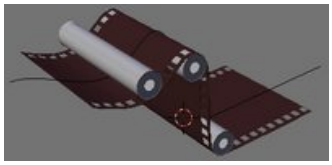
Curve modifier panel.

Ob
: The name of the curve object that will affect the deformed object.

VGroup
: A vertex group name within the deformed object. The modifier will only affect vertices assigned to this group.

X, Y, Z, -X, -Y, -Z
: This is the axis that the curve deforms along.

Example



Curve modifier example.
This example shows a mesh, the film base, that is deformed by a Curve modifier. The curve is the single black line running through the wheels. Animating the film going through the wheels is just a matter of dragging the film mesh along the curve axis.

Hooks

Hooks are similar to shape keys in that they deform a mesh over time (frames). The difference is that hooks make it look like the mesh is snagged with a fish hook. Moving the hook moves selected vertices under the influence of the hook (which is usually just an empty â← but might be any object), and you make the hook move by animating the motion of the empty through Ipo keys or any other animation method. As the hook moves, it weighted-pulls vertices from the mesh with it. If you have used proportional editing, you can think of it as animated proportional editing. While hooks do not give you the fine control over vertices movement that shape keys do, they are much simpler to use.

Object Hooks

Mode: Object mode or Edit mode

Panel: Modifiers (Editing context, F9)

Hotkey: CtrlH

Menu: Mesh Â» Vertices Â» Add Hook

Description

Hooks give access at object level to the underlying geometry of meshes, curves, surfaces or lattices. A hook is an object feature and it is like a parent to an object, but for vertices. You can create as many hooks to an object as you like, and assign for each hook vertices that will be affected. Overlapping hooks is also possible, here a weighting factor per hook is provided that determines the amount each hook will affect the overlapping vertices.

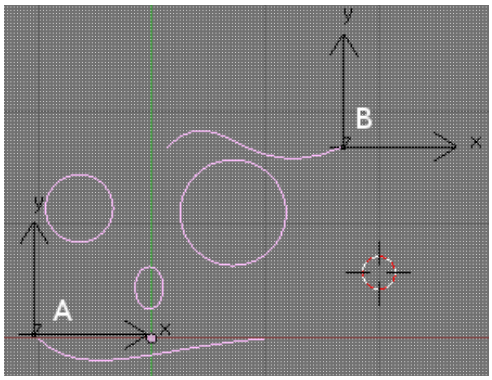
All object level options and transformations are possible on the hook object, including using hierarchies, constraints, Ipo and path animations. You can also make the hook-parent a child of the original object if you donâ← t want object transformations to deform the hooks.

Note
When you change topology (i.e. more destructive editing than just manipulating existing points), you most likely have to reassign existing hooks as well.

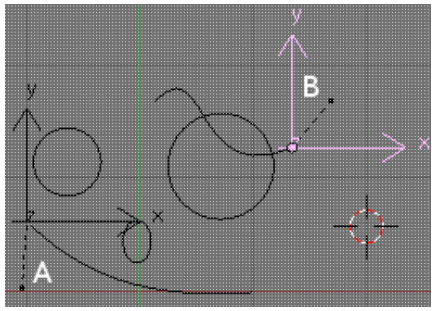
Examples

A typical example of using hooks is to animate vertices or groups of vertices. For example, you may want to animate the vertices associated with a â← Mouthâ← on a characterâ← s face.

In (*Animated face frame 1*) and (*Animated face frame 10*) a face made from BÃ©zier curves has two hooks applied. One applied to a control-point on the mouth labelled â← Aâ← and one applied to the eyebrow labelled â← Bâ← . The animation has 10 frames over which hook A moves up and hook B moves down.



Animated face frame 1.



Animated face frame 10.

Adding Hooks

Mode: Edit mode

Panel: Modifiers (Editing context, F9)

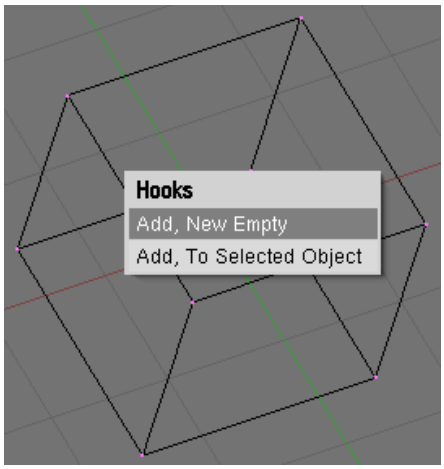
Hotkey: CtrlH

Menu: Mesh Â» Vertices Â» Add Hook

Description

Since hooks relate to vertices or control points, most editing options are available in Edit mode for meshes, curves, surfaces and lattices (not texts, as you might have guessed!).

Options



Hooks menu.

Add, New Empty
: Adds a new hook and creates a new empty object, which will be the hook controller, at the center of the selection.

Add, To Selected Object
: When another object is selected (you can do this in Edit mode with CtrlRMB), the new Hook modifier is created with that selected object as hook controller.

Editing Hooks

Mode: Edit mode

Panel: Modifiers (Editing context, F9)

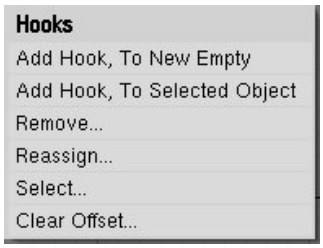
Hotkey: CtrlH

Menu: Mesh Â» Vertices Â» Add Hook

Description

Once hooks are available in an object, the Hooks menu will give additional options:

Options



Extended Hooks menu.

Remove...
: This will give a new menu with a list of hooks to remove.

Reassign...
: Use this if you want to assign new vertices to a hook.

Select...
: Select the vertices attached to a specific hook

Clear Offset...
: Neutralize the current transformation of a hook parent.

Hook Modifier

Mode: Object mode or Edit mode

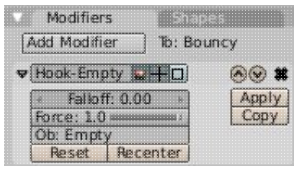
Panel: Modifiers (Editing context, F9)

Hotkey: CtrlH

Description

Hooks are modifiers, that are added to the modifier stack. For each Hook modifier, you can give a hook a new name (the default name is the hook controller name, prefixed with â← `Hook-`â←), give it a new controller, or assign it a Force weighting factor.

Options



In the Editing context (F9), Modifiers panel, when a hook is created, you can control it via this panel.

Ob
: The parent object name for the hook. Changing this name also recalculates and clears offset.

Reset
: Recalculate and clear the offset transform of hook.
The same as the Clear Offset... option from Hooks menu (CtrlH).

Recenter
: Set hook center to cursor position.

Select â← Edit mode only
: Select affected vertices on mesh.
The same as the Select... option from Hooks menu (CtrlH).

Reassign â← Edit mode only
: Reassigns selected vertices to this hook.
The same as the Reassign... option from Hooks menu (CtrlH).

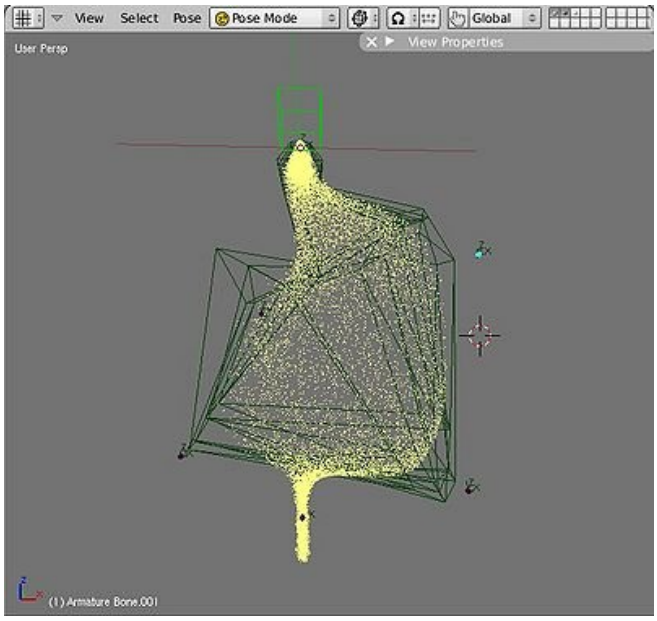
Force
: Since multiple hooks can work on the same vertices, you can weight the influence of a hook this way.
Weighting rules are:
 - ◊ If the total of all forces is smaller than **1.0**, the remainder (`1.0 - (forces sum)`, will be the factor the original position have as force.
 - ◊ If the total of all forces is larger than **1.0**, it only uses the hook transformations, averaged by their weights.

Falloff

If not zero, the falloff is the distance where the influence of a hook goes to zero. It currently uses a smooth interpolation, comparable to the proportional editing tool.

Examples

Hook modifier on lattices



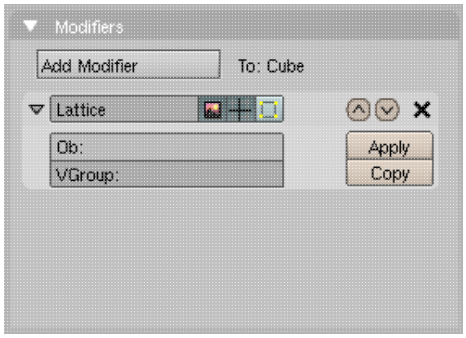
Lattice Modifier

Mode: Object mode

Panel: Modifiers (Editing context, F9)

The Lattice modifier deforms the base object according to the shape of a Lattice object.

Options



Modifier panel with Lattice modifier activated.

Ob
: The Lattice object with which to deform the base object.

VGroup
: An optional vertex group name which controls the strength of the deformation.

Hints

The Lattice modifier has the same functionality as its earlier counterpart. The primary differences between the earlier version and the implementation as a modifier is that now the effect can be applied in realtime, while editing, to **any** object type without concern for a parent/child relationship.

Instead, one only has to type the name of the lattice object into the Ob field and the modifier takes effect immediately. You can even see the effect on the object inside Edit mode using the Cage mode button (**6**).

There is a separate Lattice panel for controlling the Lattice object attributes, in the same Editing context (F9).

A lattice consists of a non-renderable three-dimensional grid of vertices. Their main use is to give extra deformation capabilities to the underlying object they control (either *via* a modifier, or being its parent). These â← victimâ← objects can be meshes, surfaces and even particles.

Why would you use a lattice to deform a mesh instead of deforming the mesh itself in Edit mode? There are a couple of reasons for that:

- First of all: itâ← s easier. Since your mesh could have a zillion vertices, scaling, grabbing and moving them could be a hard task. Instead, if you use a nice simple lattice your job is simplified to move just a couple of vertices.

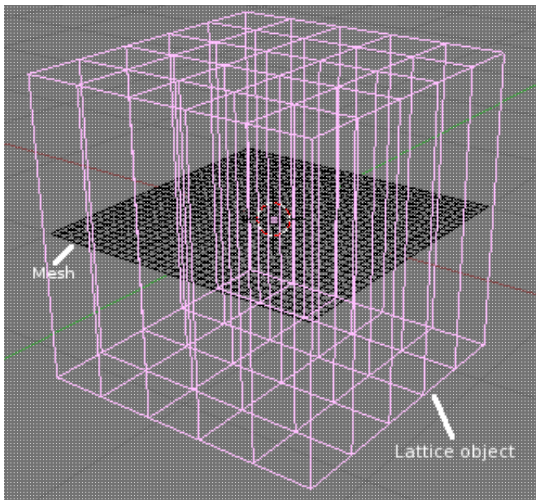
- Itâ← s nicer. The deformation you get looks a lot better!
- Itâ← s fast! You can use the same lattice to deform several meshes. Just give each object a lattice modifier, all pointing to the same lattice.
- Itâ← s a good practice. A lattice can be used to get different versions of a mesh with minimal extra work and consumption of resources. This leads to an optimal scene design, minimizing the amount of modeling work. A lattice does not affect the texture coordinates of a meshâ← s surface. Subtle changes to mesh objects are easily facilitated in this way, and do not change the mesh itself.

Example/Tutorial(s)

There are example tutorials in the Tutorials section. One shows how to shape a fork and the other shows how to make one object follow the shape of another.

Here is a really quick example of creating landscape terrain.

Note
I use the Top 3D view (7 NumPad) to perform some of these actions.



Lattice and plane.
First start with a plane Mesh object and, in Edit mode, subdivide it about five times using the Subdivide button in the Mesh Tools panel, and click All Edges (Mesh Tools More panel), so we can see all the faces in wireframe mode.

And then add a Lattice object to the scene. The default name for the Lattice object is â← `Lattice`â← ! Scale the Lattice to fit around the Mesh object (S).

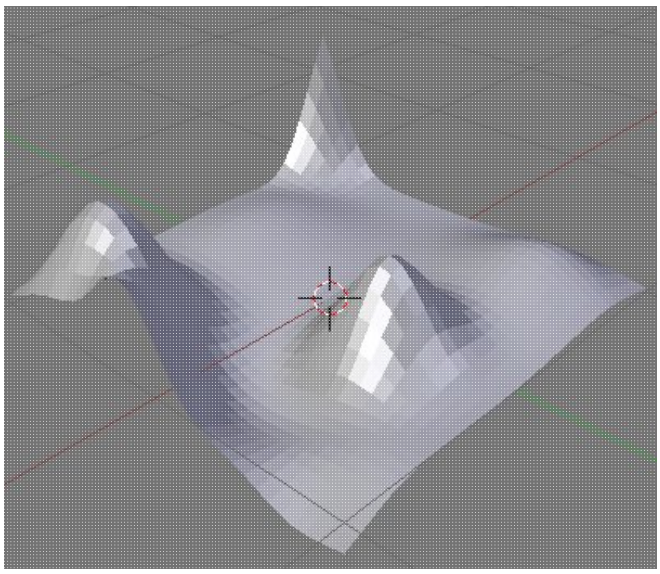
At this time the Lattice and the Mesh are independent of each other. Neither knows of the other. We associate them using a Lattice modifier. Go ahead and select the Mesh object and *add* a Lattice modifier to it. Then fill in the Ob field with the name of the Lattice object (â← `Lattice`â← by default). Be careful, the field is case sensitive.

Go back to the Lattice object and bump up the U and V subdivisions to **6** each. (*Lattice and plane*) is what we have so far as viewed from a perspective 3D view.

Now we can begin to deform the Mesh object by going into Edit mode while the Lattice object is selected. Grab a vertex, or two, and drag them (G) up the Z-axis (Z). Then grab a few more and move them down the Z-axis. Note, if you feel you donâ← t have enough vertices on the Lattice object you can always exit Edit

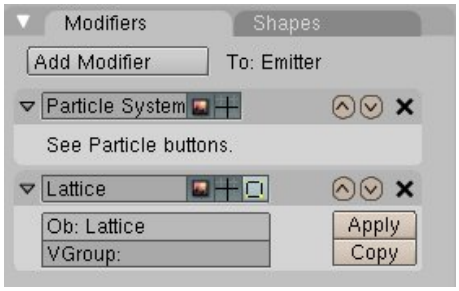
mode (i.e. go back to Object mode) and bump up the U and V subdivisions even more.

(*Quick landscape complete*) is my attempt at creating a landscape. It is in solid mode with the Lattice object hidden on another layer (M).



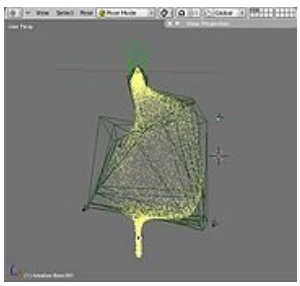
Quick landscape complete.

Particles and Lattices



Particles following a lattice.
Particles follow a Lattice if the modifier sequence is right. First the particles, then the lattice!

Example 1



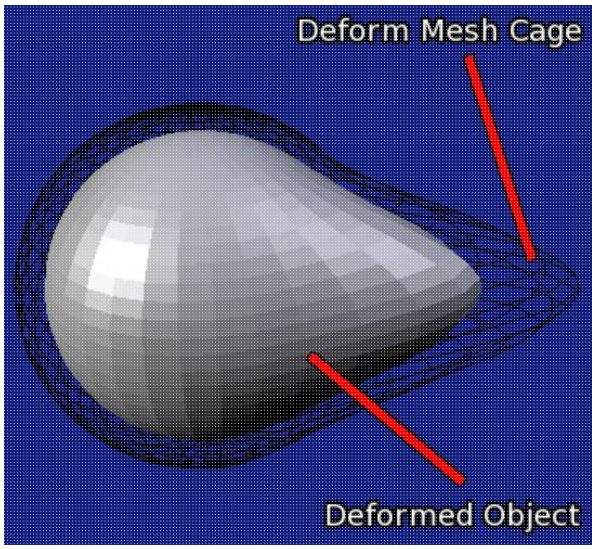
Particles flow being effected by a lattice being deformed by armature bones parented to hook modifiers on the lattice.

MeshDeform Modifier

Mode: All modes

Panel: Modifiers (Editing context, F9)

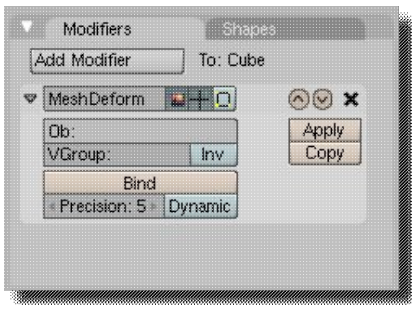
Description



Example of a deform mesh cage â← Â deforming a sphere.
The MeshDeform modifier allows an arbitrary closed mesh (of any closed shape, not just the cuboid shape of a Lattice modifier) to act as a deformation cage around another mesh.

To the right is an example of a deformed object (shaded in gray) inside the confines of a deform mesh cage (shown in black wireframe). The deform mesh cage can be any shape of mesh but it must be closed. In the example to the right it is just an UVsphere that has been altered using proportional editing, as a result the deformed object alters its shape in response. It has had a MeshDeform modifier added to it and has been told to use the deform mesh cage as the mesh it will use to deform itself.

Options



MeshDeform modifier panel.
The MeshDeform modifier is reasonably easy to use but it can be very slow to do the calculations it needs, to properly map the deform mesh cage to the deformed object.

Ob

The Ob text field is used to indicate which object the MeshDeform modifier should use for its deform mesh cage.

VGroup
: The VGroup text field is used to indicate that only the vertices in the specified vertex group will be affected by the deform mesh cage.

Inv
: Inverts the influence set by the vertex group defined in previous setting (i.e. reverts the weight values of this group).

Bind
: The Bind button is what tells the MeshDeform modifier to actually link the deform mesh cage to the deformed object, so that altering the shape of the deform mesh cage actually alters the shape of the deformed object.
Be aware that depending on the settings of the MeshDeform modifier and complexity of the deform mesh cage and/or deformed object, it can take a long time for this operation to complete. This can result in Blender not responding to userâ← s actions until it has completed, it is even possible that Blender will run out of memory and crash. As Blender progresses through this operation, you should see its progress bar (the colored bar, in the header of the User Preferences window, usually in top-right corner) progress and change of color.

Unbind
: When a deformed object has been associated to a deform mesh cage, it can later be disassociated by selecting the Unbind button which replaced the Bind one.
When Unbind is clicked, the deform mesh cage will keep its current shape, it will not reset itself back to its original start shape. If you need its original shape, you will have to save a copy of it before you alter it. The deformed object will however reset back to its original shape that it had before it was binded to the deform mesh cage.

Precision
: The Precision numeric slider field controls that accuracy with which the deform mesh cage alters the deformed object, when the points on the cage are moved.
The range of values for the Precision field can range from **2** to **10**, the default being **5**. Raising this value higher can greatly increase the time it takes the MeshDeform modifier to complete its binding calculations, but it will get more accurate cage mapping to the deformed object. This rise in calculation time can make Blender stop responding until it has calculated what it needs to. As well as making Blender not respond, raising the Precision value high and then trying to Bind on a very complex deform mesh cage and/or deformed object can use large amounts of memory and in extreme cases crash Blender. To be safe, save your blend file before proceeding!
This setting is unavailable once a cage is binded.

Dynamic
: The Dynamic button indicates to the MeshDeform modifier that it should also take into account deformations and changes to the underlying deformed object which were not a direct result of deform mesh cage alteration.
With Dynamic button activated, other mesh altering features (such as other modifiers and shape keys) are taken into account when binding a deform mesh cage to the deformed object, increasing deformation quality. It is deactivated by default to save memory and processing time when bindingâ←¦
As Precision, this setting is unavailable once a cage is binded.

Other Characteristics of MeshDeform Modifier

There are some characteristics of the MeshDeform modifier which are not directly visible within its control panel. This section list some of those issues and features.

Mode of Operation

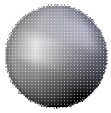
Alterations made to the deform mesh cage will only be reflected in the deformed object when the cage is in Edit mode, when in Object mode the cage can be scaled and distorted but it will not effect the deformed object.

Deform Mesh Cage Location after Binding

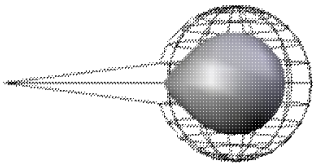
While a deform mesh cage is being binded to a deformed object, the cage must surround all the parts of the deformed object you wish to be affected by the cage. Once the deform mesh cage has been binded it can be moved away from the deformed object in Object mode. When you then switch the deform mesh cage back to Edit mode and alter its shape, it will alter the deformed object even when it is not directly surrounding it.

Distance from Deform Mesh Cage & Deformed Object

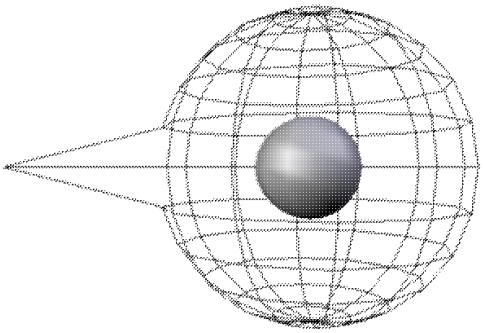
Distance between the deform mesh cage and the object to be deformed has an influence on the amount of change imparted to the deformed object when the cage is altered (when in Edit mode). When the deform mesh cage is further away from deformed object, then the amount of change imparted to the later is less and less local to a specific area of the deformed object. When the cage is closer to the deformed object, the amount of influence upon the later is greater and more local to a specific area on the deformed object.



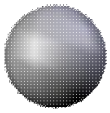
Un-deformed sphere.



Small distance between a deform mesh cage & a deformed object.



Large distance between a deform mesh cage & a deformed object.



Animation showing the difference between each sphere deform in a deform mesh cage (click to see animation).

Above are examples of the effect of different deform mesh cage distances from a deformed object. The top left image shows a normal un-deformed UVsphere. The top right image shows the same UVsphere but with a deform mesh cage which is very close to the deformed object, and as a result there is quiet large deformation in the deformed object. The bottom left image shows the deformation of a deformed object when the deform mesh cage is further away. It can be seen that the deformed object alteration is much more muted, even though the vertex that has been moved in the deform mesh cage has moved by the same amount. The bottom right image shows an animated version of the other three images showing the change in the deformed object with different deform mesh cage distances.

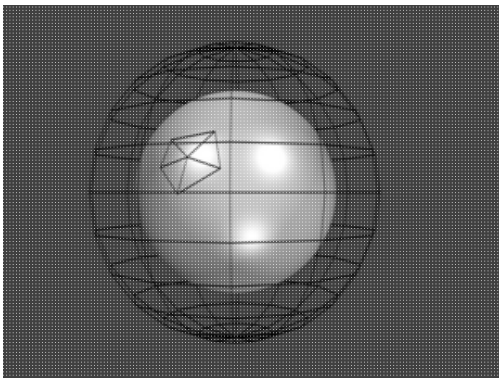
Multires support

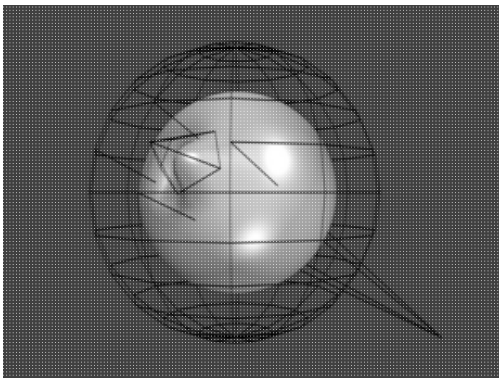
Blender 2.49 can now handle both MeshDeform modifier and Multires feature â← Â with the limitation that MeshDeform will only work when the multires level it was binded with is active. This means that when you are not working on the multires level you used during binding, the modifier (and hence any cage modification) will have no effect on the deformed object. Of course, you can add as much MeshDeform modifiers as you have multires levels.

Note however that the MeshDeform modifier does not seem to take advantage of increased level of details produced by higher multires levelsâ←¦

Interior Control

Besides the cage itself, you can have more faces within the cage, which can form a â← sub-meshâ← loose or closed, and allowing an extra control on some areas of the deformed object. These sub-meshes might be linked or not to the main cage, intersect it, etc. For example, you could use a big cage deforming a whole face (to get â← cartoonâ← effects), with two smaller sub-meshes around the eyes, to better control them.







Non-deformed Sphere, with a deforming cage including a small sub-mesh.



Deformed Sphere in two different directions by the two sub-meshes of the cage.

In the example above, much more modest, we are back with our two spheres, the small one deformed by the big one. But this time, a small open sub-mesh (a slightly concave plane) has been added to the deform mesh cage, closer from the deformed object (left picture). In the right picture, the sphere-cageâ← s vertex has been moved towards the center, causing a quite general recess of the deformed sphere, since the cage is originally quite far away from it. However, the area of the deformed object controlled by the little added sub-mesh (of which the central vertex has been moved in the opposite direction) comes out as a little central bump.

MeshDeform Characteristics
The MeshDeform modifier works by taking into account bounding faces of the cage to the deformed mesh when doing itâ← s calculations to bind the cage to the deformed object, influencing the implicit object (e.g. vertices and edges implicitly), using their boundary faces to do itâ← s bind calculations. (Info from Greylica)

Implementation

The MeshDeform modifier implementation method (in Blender 2.46) is currently â← Harmonic Coordinates (for Character Articulation)â← Volume Deformation Method developed by Pushkar Joshi, Mark Meyer, Tony DeRose, Brian Green and Tom Sanocki of Pixar Animation Studios. This method was presented at the Siggraph 2007 conference. It has many advantages in controlling the deformations of meshes.

A copy of the implementation PDF document can be downloaded at graphics.pixar.com.

A video demonstrating some of the important features of the implementation can be viewed below:

Examples

- Blender file showing an example of a MeshDeform modifier.
- Blender file showing the effect of distance on the effect of a MeshDeform modifier.
- Blender file showing the usage of deform mesh cage on the BBB Chinchilla.

See Also

- The Lattice modifier.

Shrinkwrap Modifier

Mode: All modes

Panel: Modifiers (Editing context, F9)

Description

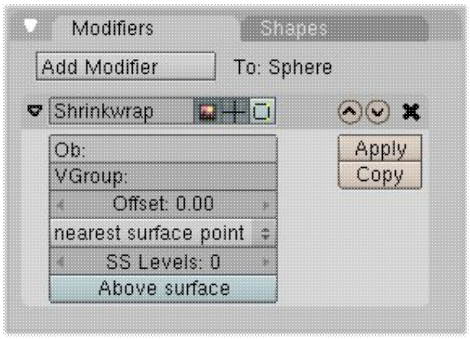
The Shrinkwrap modifier allows an object to â← shrinkâ← to the surface of another object. It moves each vertex of the object being modified to the closest position on the surface of the given mesh (using one of the three methods available). It can be applied to meshes, lattices, curves, surfaces and texts.

As for most of the deform modifiers, the affected â← verticesâ← are the â← computedâ← one, i.e. the real geometry of the object at the time the modifier is calculated, and not the original *vertices*/control points.

In some way, Shrinkwrap acts like a view-independent retopo tool, since the vertices are projected along their normal or moved to the nearest surface point. But it doesnâ← t give accuracy problems as retopo, since it works in object space instead of image space. Also itâ← s possible to â← keep a distanceâ← from the target position.

Note
For those who found the Shrinkwrap modifier pretty useful, but would like it to move empties or objects positionsâ←¦ Have a look at the Shrinkwrap constraint!

Options



The default Shrinkwrap modifier panel.

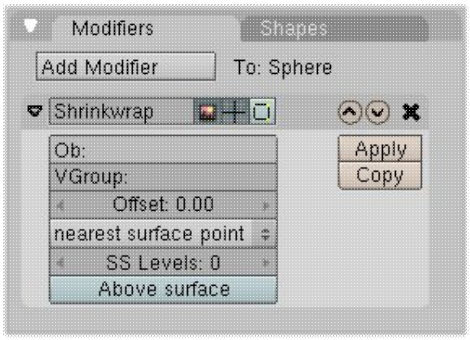
Ob
: Shrink target, the mesh to shrink/wrap around.

VGroup
: The weight paint for this vertex group of the current modified mesh will controls whether and how much each vertex is displaced to its target position. If a vertex is not a member of this group, it is not displaced (same as weight 0)

Offset
: The distance that must be kept from the calculated target position, in Blender Units.

SS Levels
: This should apply a (temporary) Catmull-Clark subsurf to the modified object (or itâ← s target?), before computing the wrapâ←¦ But it does not seem to have any effectâ←¦

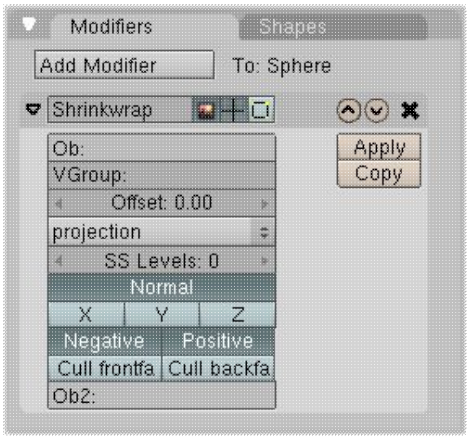
Shrinkwrap type
: This drop-down list specifies the method to be used to determine the nearest point on the targetâ← s surface for each vertex of the modified object. Some options will add some extra, specific controls to the panel.



The Shrinkwrap modifier with nearest surface point method.

nearest surface point
: This will select the nearest point over the surface of the shrink target. It adds one extra option

Above surface
: When enabled, the shrunk vertices are always kept above (on the front side) of their â← floor faceâ← (i.e. the face they are the nearest from). This is only meaningful when Offset is not null.



The Shrinkwrap modifier with projection method.

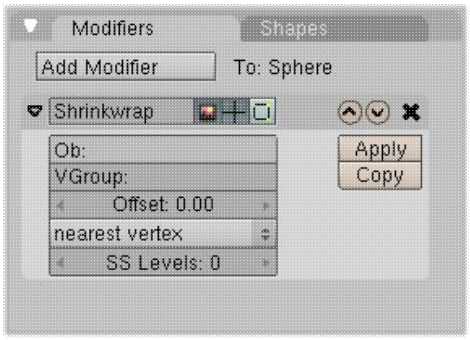
projection
: This will project vertices in a given direction, until they touch the shrink target. *The vertices never touching the shrink target are left in their original position*. This implies that, depending on the settings of this option (especially when using the Normal projection axis) and the relative positions of the two objects, the modified object might sometimes remain undeformedâ←¦ This is not a bug, just â← playâ← with the settings (especially the Negative/Positive ones), or move one of the objects aroundâ←¦
This method is the hardest to master, as it might sometimes gives unexpected resultsâ←¦ It adds quite a few extra options:

Normal, X, Y, Z
: Along which direction the projection is done. Normal will use the modified objectâ← s normals as project â← axisâ← , and is mutually exclusive with the other three options. X, Y and Z will use the corresponding local axis of the modified object â← Â they can be combined with each other, yielding a â← median axisâ← of projection.

Negative and Positive
: This allows you to select the allowed direction(s) of the shrink along the selected axis.

Cull frontfaces, Cull backfaces
: This allows you to prevent any projection over the â← front sideâ← (respectively the â← back sideâ←) of the targetâ← s faces. The â← sideâ← of a face is determined by its normal (front being the side â← from whereâ← the normal â← originatesâ←).

0b2
: An additional object to project overâ←¦



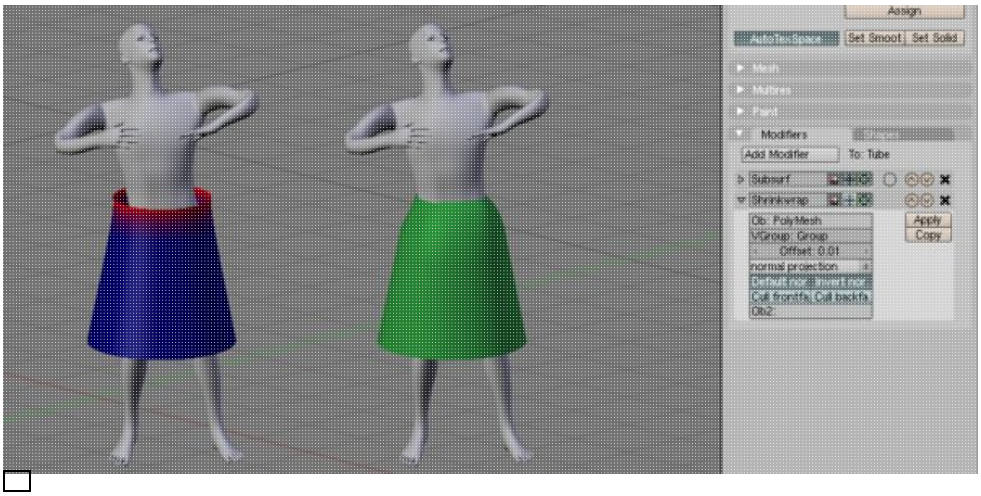
The Shrinkwrap modifier with nearest vertex method.

nearest vertex
: This will select the nearest vertex of the shrink target. It adds no extra option.

Examples

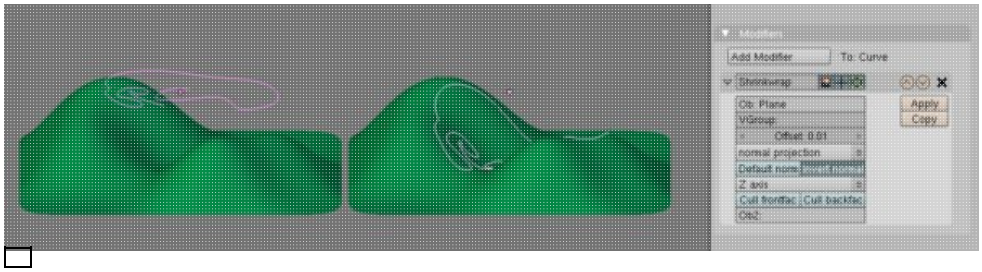
Artists have found many uses for this modifier. Here are some of those showing Shrinkwrap modifier utility.

Note
Although most showed examples are directed to cloth modeling, this modifier is very useful in the modeling process (as for example doing retopology, skin bulgingâ←¦).

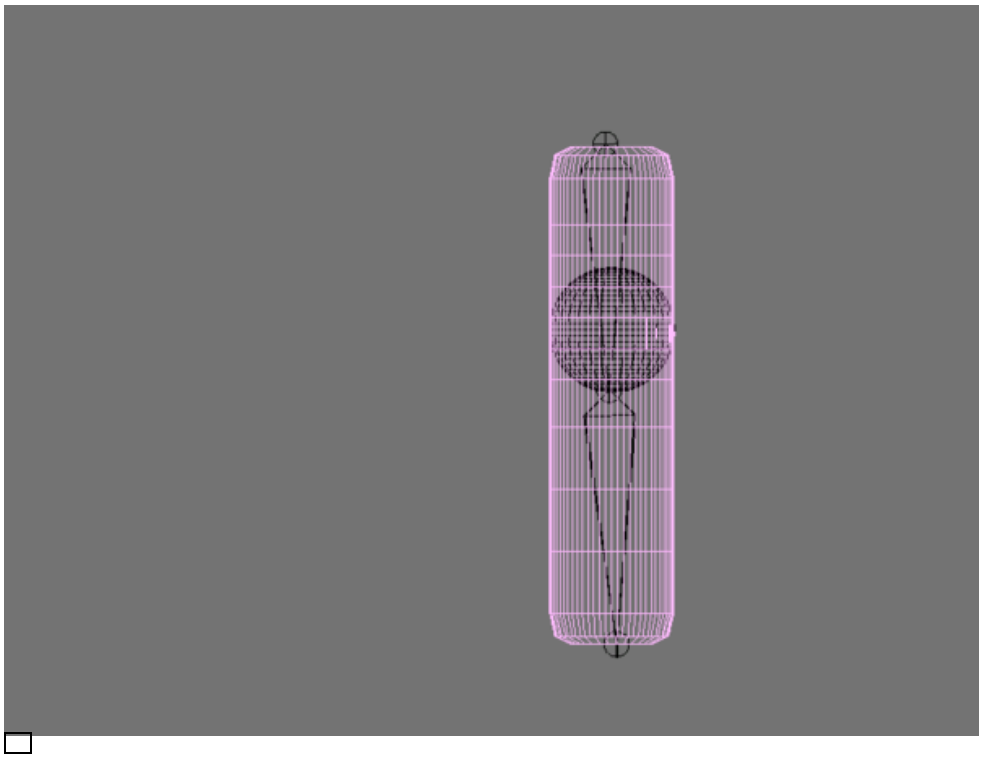


Shrinkwrapping a skirt to a model: simply paint the vertices that are supposed to be close to the targetâ ← s mesh.

Basic example of shirt making: just make a very simple cage and shrinkwrap it to the model.



Shrinkwrapping a curve to a terrain: apply a projection along the Z axis.



Using shrinkwrap to achieve skin slide (Renderdemon idea).

SimpleDeform Modifier

Mode: All modes

Panel: Modifiers (Editing context, F9)

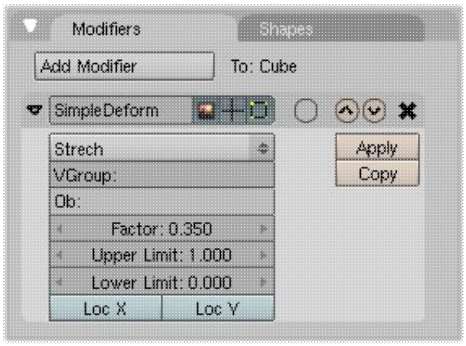
Description

The SimpleDeform modifier allows to easily apply some simple deform functions to objects (meshes, lattices, curves, surfaces and texts are supported).

As for most of the deform modifiers, the deform functions are applied to the â← computed verticesâ← , i.e. to the real geometry of the object at the time the modifier is calculated, and not to the original *vertices*/control points. This means you can increase the level of detail of the deform effect by inserting first a Subsurf modifier (for meshes), or raising the Resol U and V settings (for curves/surfaces/texts).

Using another object, itâ← s possible to define the axis and origin of the deformation, allowing to apply very different effects.

Options

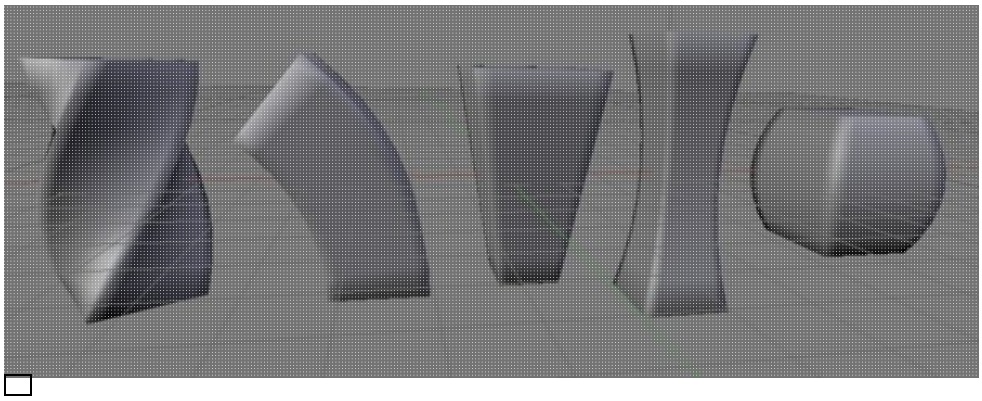


The SimpleDeform modifier panel.

Deform type

This drop-down list defines the deform function applied, among four available:

- ◊ Twist â← Rotates around the Z axis.
- ◊ Bend â← Bends the mesh over the Z axis.
- ◊ Taper â← Linearly scales along Z axis.
- ◊ Stretch â← Stretches the object along the Z axis (negative Factor leads to squash).



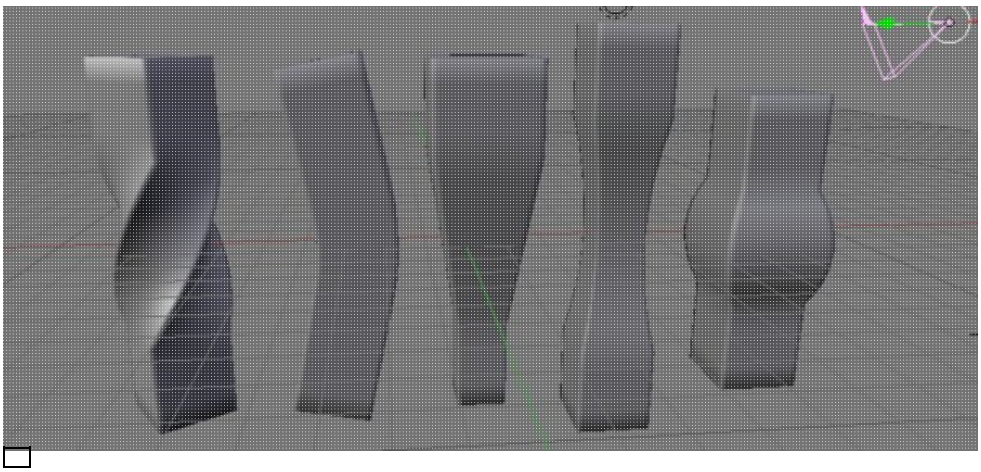
The different deform functions, from left to right:
Twist, Bend, Taper, Stretch with positive Factor, Stretch with negative Factor.

VGroup
: The name of the vertex group that indicates whether and how much each vertex is influenced by the deformation.

Ob
: An object that defines the origin of deformation (usually an empty). This object can be:
 ◊ Rotated to control axis (as it is its Z axis that is now used as â← guideâ←).
 ◊ Translated to control the origin of deformation.
 ◊ Scaled to change the deform factor.

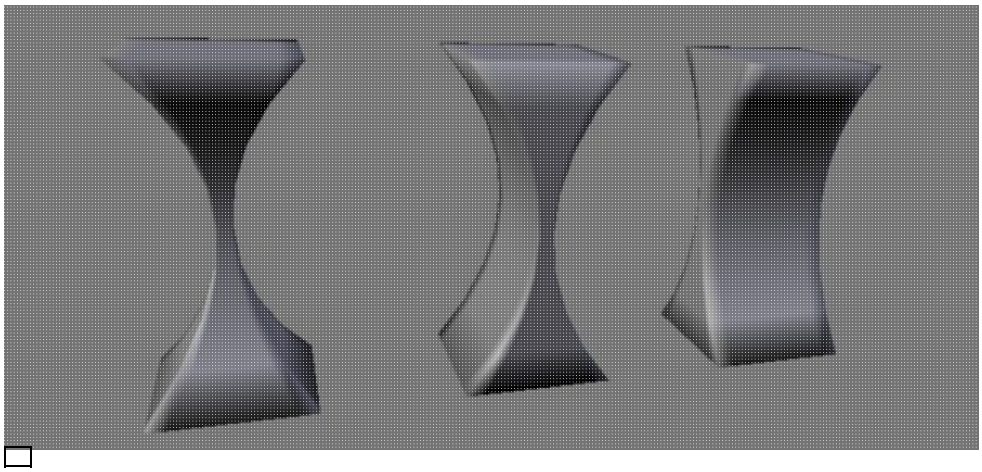
Factor
: The amount of deformation.

Upper Limit and Lower Limit
: These settings allow you to limit the part of the object that is affected by the deformation (they are proportional values, from **0.0** to **1.0**). Obviously, Upper Limit canâ← t be lower than Lower Limitâ←¦



The various deform functions with Upper Limit set to **0.75** and Lower Limit to **0.25**.

Lock X/Loc Y (Taper and Strecth types only)
: This controls whether the X and/or Y coordinates are allowed to change or not. Thus it is possible to squash the X coordinates of an object and keep the Y coordinates intactâ←¦



A Stretch SimpleDeform modifier without lock, with Loc X, and with Loc Y.

Notes

When the object controlling the origin (the one in the Ob field) is a child of the deformed object, this creates a cyclic dependency in Blenderâ← s data system (the DAG â← Â â← dependency graphâ← ?). The work around is to create an empty and attach both objects to itâ←¦

Wave Modifier

Mode: Object mode

Panel: Modifiers (Editing context, F9)

Description

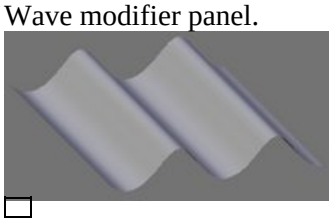
The Wave modifier adds an ocean-like motion to the Z coordinate of the objectâ ← s vertices/control points. This modifier is available for meshes, lattices, curves, surfaces and texts, with a few restrictions for non-*mesh* objects:

- Activating Normals or typing a name in VGroup will then simply deactivate the modifier.
- Even worth, selecting UV as texture coordinates will make Blender crash at once!

Options



Wave modifier panel.



Linear wave front.



Circular wave front.

This panel is found in the Editing context (F9), Modifiers panel, selecting Wave as the modifier.

X and Y
: The wave effect deforms vertices/control points in the Z direction, originating from the given starting point and propagating along the object with circular wave fronts (both X and Y activated), or with rectilinear wave fronts, then parallel to the axis corresponding to the X or Y button activated.

Cycl
: Repeats the waves cyclically, rather than a single pulse.

Normal â← Meshes only!
: Displaces the mesh along the surface normals (instead of the objectâ← s Z-axis).

Time sta
: The frame at which the wave begins (if Speed is positive), or ends (if Speed is negative). Use a negative frame number to prime and pre-start the waves.

Lifetime
: The number of frames in which the effect lasts. **0** means foreverâ←¦

Damptime
: An additional number of frames in which the wave slowly dampens from the Height value to zero after Lifetime is reached. The dampening occurs for all the ripples and begins in the first frame after the Lifetime is over. Ripples disappear over Damptime frames.

Sta X and Sta Y
: The starting point of the waves, in the objectâ← s local coordinates.

Ob
: Use another object as reference for starting position of the wave. Leave blank to disable. Note that you then can animate this objectâ← s position, to change the waves origin along timeâ←¦

VGroup â← Meshes only!
: A vertex group name, used to control the parts of the mesh affected by the wave effect, and to what extent (using vertex weights).

Texture
: Use this texture to control the objectâ← s displacement level. Animated textures can give very interesting results hereâ←¦

Texture Coordinates
: This menu lets you choose the textureâ← s coordinates for displacement, between Local, Global, Object (which adds an additional Ob field just below, to type in the name of the object from which to get the texture coordinates), and UV (which adds an extra UV Layer drop-down list, where to select the UV layer to be used â← Â **warning:** do not activate this options with non-mesh objects, it seems to make Blender crash at once!).

Speed
: The speed, in BU (for â← Blender Unitsâ←) per frame, of the ripple.

Height
: The height or amplitude, in BU, of the ripple.

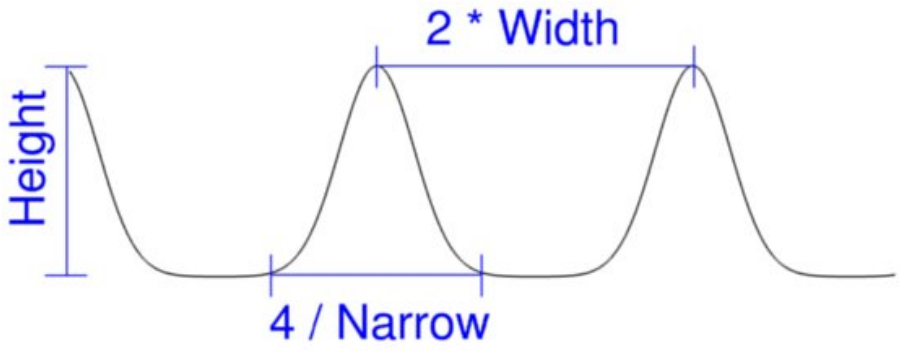
Width
: Half of the width, in BU, between the tops of two subsequent ripples (if Cycl is enabled). This has an indirect effect on the ripple amplitude â← Â if the pulses are too near to each other, the wave may not reach the **0** Z-position, so in this case Blender actually lowers the whole wave so that the minimum is zero and, consequently, the maximum is lower than the expected amplitude. See technical details below.

Narrow
: The actual width of each pulse, the higher the value the narrower the pulse. The actual width of the area in which the single pulse is apparent is given by `4/Narrow`. That is, if Narrow is **1** the pulse is **4** units wide, and if Narrow is **4** the pulse is **1** unit wide.

Warning
All the values described above must be multiplied with the corresponding Scale values of the object to get the real dimensions. For example, if the value of Scale Z is **2** and the value of Height of the waves is **1**, it gives us final waves with a height of **2BU**!

Technical Details and Hints

The relationship of the above values is described here:



Wave front characteristics.

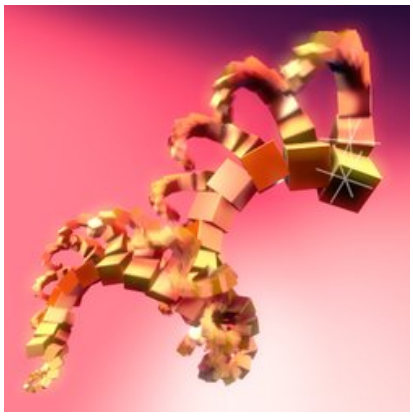
To obtain a nice wave effect similar to sea waves and close to a sinusoidal wave, make the distance between following ripples and the ripple width equal, that is the Narrow value must be equal to `2/Width`. E.g. for Width=**1**, set Narrow to **2**.

Array Modifier

Mode: Any mode

Panel: Modifiers (Editing context, F9)

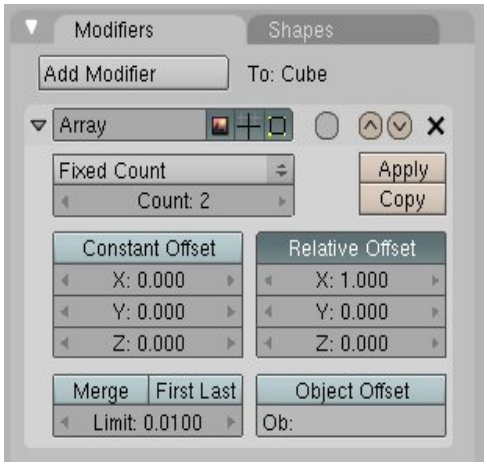
Description



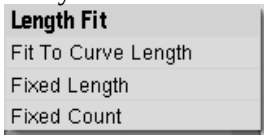
Multidimensional array animated with motion blur.
The Array modifier creates an array of copies of the base object, with each copy being offset from the previous one in a number of possible ways. Vertices in adjacent copies can be merged based on a merge distance, allowing smooth subsurf frameworks to be generated.

This modifier can be useful when combined with tilable meshes for quickly developing large scenes. It is also useful for creating complex repetitive shapes.

Options



Array modifier.



Length Fit menu.

Length Fit menu

Controls how the length of the array is determined, see (*Length Fit menu*). There are three choices, activating respectively the display of the Ob, Length or Count setting:

- ◊ Fit To Curve Length â← Generates enough copies to fit within the length of the curve object specified in Ob.
- ◊ Fixed Length â← Generates enough copies to fit within the fixed length given by Length.
- ◊ Fixed Count â← Generates the number of copies specified in Count.

Ob

The Curve object to use for Fit To Curve Length.

Length

The length to use for Fixed Length.

Count

The number of duplicates to use for Fixed Count.

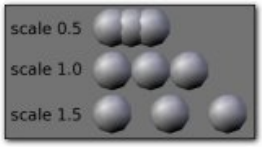
Notes

- Both Fit To Curve Length and Fixed Length use the local coordinate system size of the base object, which means that scaling the base object in Object mode will not change the number of copies generated by the Array modifier. Applying the scale (CtrlA) can be useful in this case.
- Fit To Curve Length uses the local coordinate system length of the curve, which means that scaling the curve in Object mode will not change the number of copies generated by the Array modifier. Applying the scale (CtrlA) can also be useful in this case.

Constant Offset, X, Y, Z

Adds a constant translation component to the duplicate objectâ← s offset. X, Y and Z constant components can be specified.

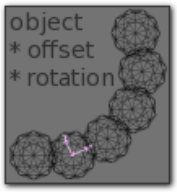
Relative Offset, X, Y, Z



Relative offset example.

Adds a translation equal to the objectâ← s bounding box size along each axis, multiplied by a scaling factor, to the offset. X, Y and Z scaling factors can be specified. See (*Relative offset example*).

Object Offset, Ob



Object offset example.

Adds a transformation taken from an object (relative to the current object) to the offset. See (*Object offset example*).

Merge
: If enabled, vertices in each copy will be merged with vertices in the next copy that are within the given Limit merge distance.

First Last
: If enabled **and** Merge is enabled, vertices in the first copy will be merged with vertices in the last copy (this is useful for circular objects, see (*First Last merge example*)).



Subsurf discontinuity caused by not merging vertices between first and last copies (First Last off).



Subsurf discontinuity eliminated by merging vertices between first and last copies (First Last on).

First Last merge example.

Limit
: Controls the merge distance for Merge.

Start cap
: The mesh object to be used as a start cap. A single copy of this object will be placed at the â← beginningâ← of the array â← Â in fact, as if it was in position **-1**, i.e. one â← array stepâ← before the first â← regularâ← array copy. Of course, if Merge is activated, and the Start cap is near enough of the first copy, they will be merged.

End cap
: The mesh object to be used as an end cap. A single copy of this object will be placed at the â← endâ← of the array â← Â in fact, as if it was in position **n+1**, i.e. one â← array stepâ← after the last â← regularâ← array copy. And as Start cap, it can be merged with the last copyâ←¦

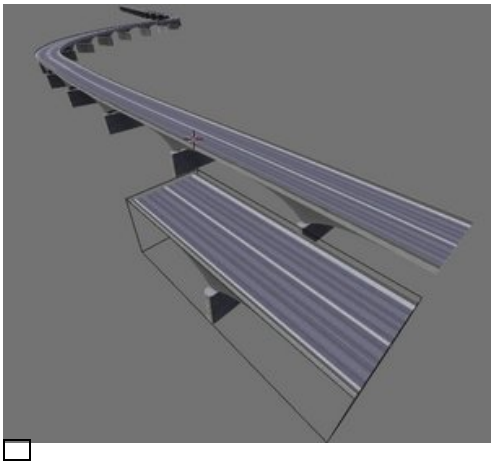
Hints

Offset Calculation

The transformation applied from one copy to the next is calculated as the sum of the three different components (Relative, Constant and Object), all of which can be enabled/disabled independently of the others. This allows, for example, a relative offset of **(1, 0, 0)** and a constant offset of **(0.1, 0, 0)**, giving an array of objects neatly spaced along the X axis with a constant **0.1BU** (Blender Units) between them, whatever being the original objectâ← s size.

Examples

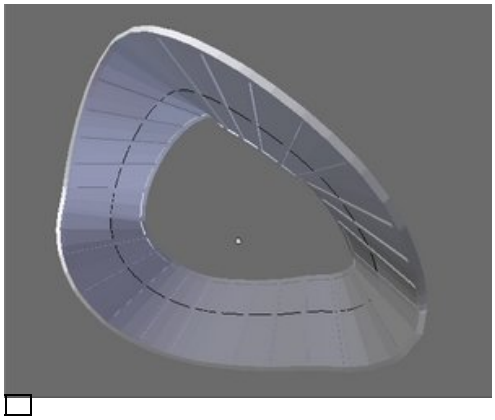
Mechanical



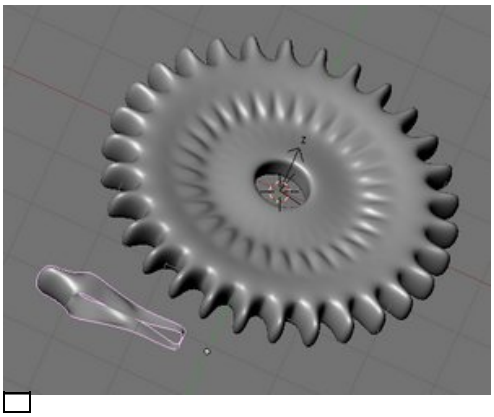
A bridge made from a tilable mesh.

Note

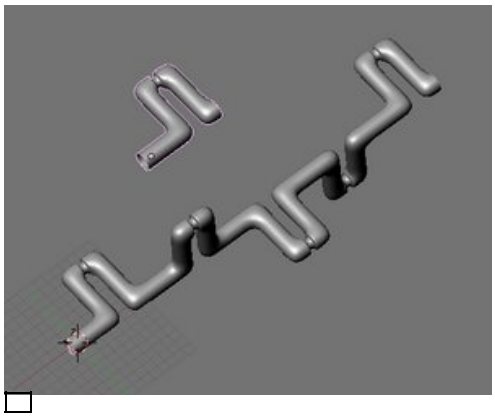
As the Curve modifier could not be after Array in the modifier stack (at the time this image was created), the Array modifier was applied (i.e. the Apply button was pressed) before the curve was added in the bridge image.



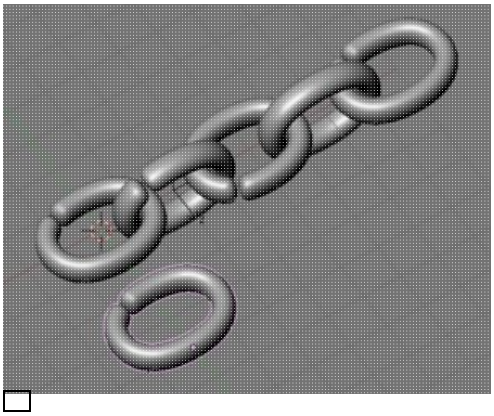
A track.
Sample blend file



A cog created from a single segment.
Sample blend file



A crankshaft.
Sample blend file

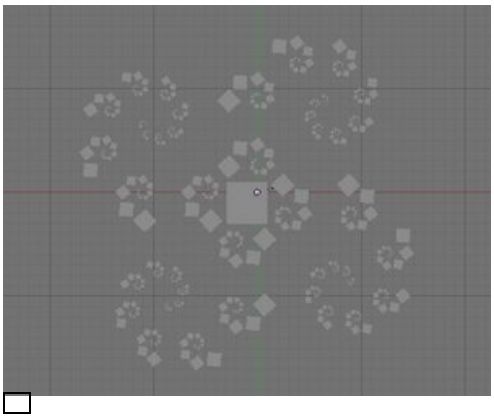


A chain created from a single link.
Sample blend file

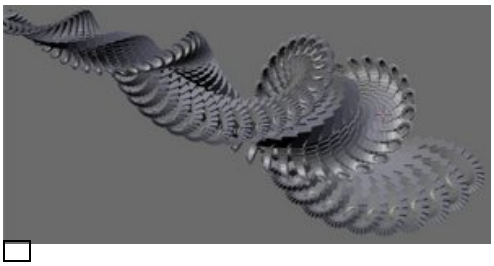
Fractal



Multidimensional array animated with motion blur.



A fractal-like image created with multiple array modifiers applied to a cube.
Sample blend file

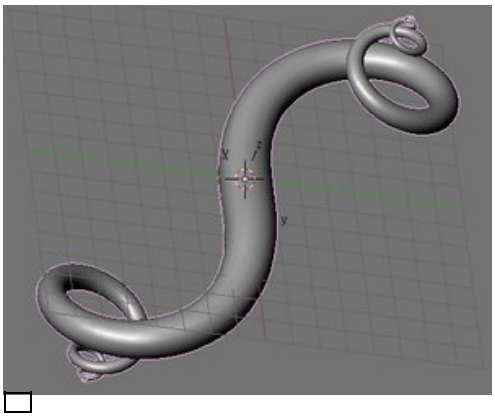


A fractal fern image created with 2 array modifiers and 1 mirror applied to a cube.

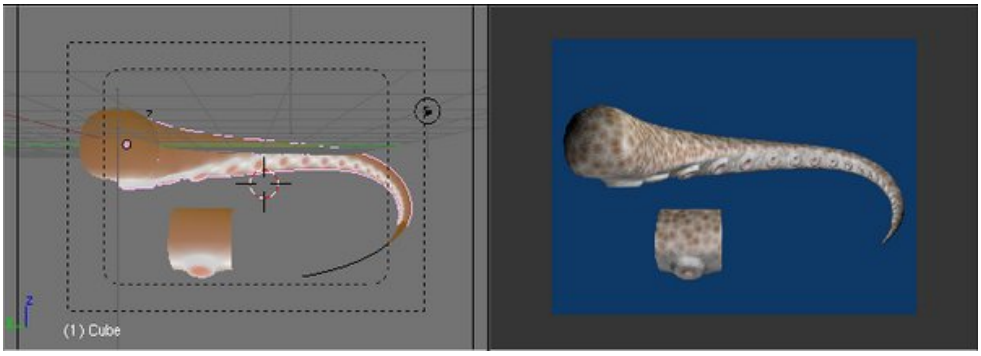
Organic



Subsurfed cube array with 1 object offset, 4 cubes and a high vertex merge setting to give the effect of skinning.



A double spiral created with two array modifiers and one subsurf modifier applied to a cube. As above, the vertex merge threshold is set very high to give the effect of skinning.
Sample blend file





A tentacle created with an Array modifier followed by a Curve modifier. The segment in the foreground is the base mesh for the tentacle; the tentacle is capped by two specially modeled objects deformed by the same Curve object as the main part of the tentacle.
Sample blend file

Tutorials

Some tutorials that exploit the Array modifier:

- Creating A Double Helix With Modifiers

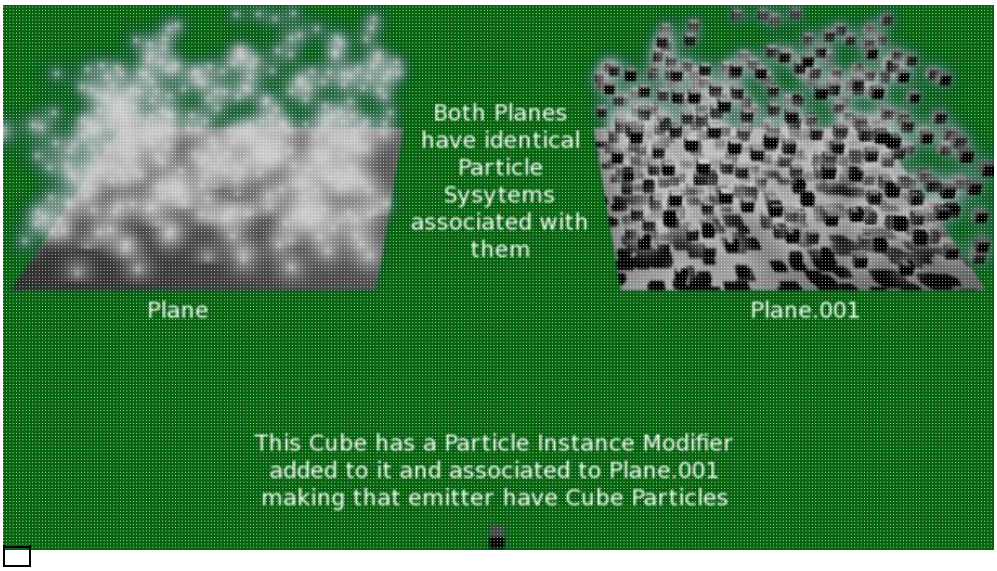
ParticleInstance Modifier

Mode: Any mode

Panel: Modifiers (Editing context, F9)

Description

When a ParticleInstance modifier is added to an object, that object will be used as a particle shape on an object which has a particle system associated with it. This means that to use this modifier you must also have another object which has a particle system on it, otherwise the ParticleInstance modifier will appear to do nothing.



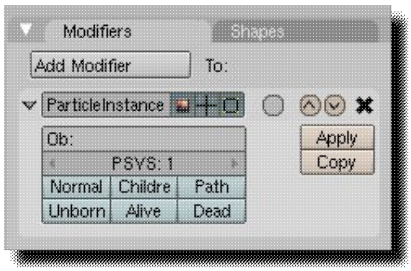
Particle system on left has no ParticleInstance modified object associated with it. The one on the right is associated with cube shown by using a ParticleInstance modifier on the cube.

Definitions/Terms

Here is a brief explanation of the various terms and definition used in relation to particles and the ParticleInstance modifier:

- Particle system â← An object (mesh) which has the ability to emit/generate particles activated on it.
- Normal particle â← A particle that is not a children/child generated particle type.
- Children/child particle â← A particle type that is generated and placed with relation to other normal particles that already exist. Children particles are generally much quicker to calculate.
- Unborn particle â← A particle which has not yet been displayed/emitted because it is not its time to be emitted/displayed. One of the reasons a particle can be in unborn state is that it is before the frame at which it is to be emitted.
- Alive particle â← A particle which has been displayed/emitted and has not yet reached its dead state. One of the reasons a particle can be in an alive state is that it has been alive for less frames than its life length.
- Dead particle â← A particle which has been displayed/emitted and has reached its end of life length and at that point it enters the dead state.

Options



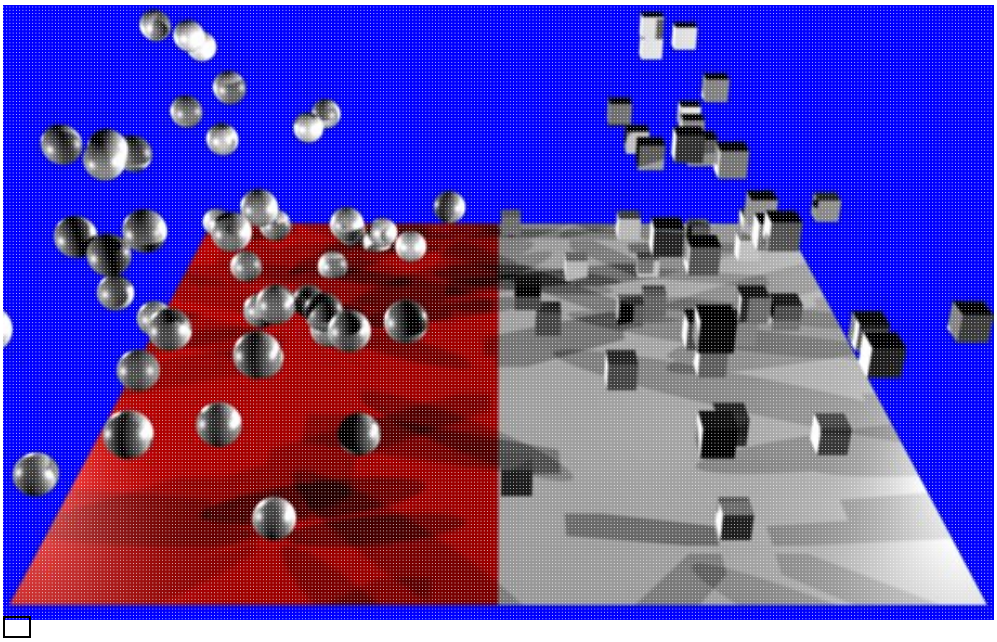
ParticleInstance modifier panel.
Because of the co-dependant way in which the ParticleInstance modifier is influenced by the underlying particle systems on other objects, some of the apparent effects generated by the ParticleInstance modifier can look and act vastly different, depending on the underlying settings of the particle systems it is associated with. This is worth taking account of if the ParticleInstance modifier settings donâ← t appear to be giving the results expected, as it may indicate that the particle system settings may need altering rather than the ParticleInstance modifier settings.

Ob

The Ob field (short of â← Objectâ←), associates this ParticleInstance modifier with another object (usually an object having a particle systemâ←¦). This indicates that when the object named in this field emits particles, those particles will have the mesh shape of the current ParticleInstance modifierâ← s mesh.
If for example a sphere has a ParticleInstance modifier added to it, when the Ob field of this modifier is filled in with the name of an object that emits particles, those particle will be sphere shaped.
Even though most of the time the Ob field will have the name of an object with a particle system, this is not mandatory, you can enter an objectâ← s name which does not have a particle system, and it will be accepted by the Ob field, as there do not appear to be any checks made to make sure the objectâ← s name entered into this field is â← validâ← .

PSYS

The PSYS field (short of â← Particle SYStemâ←), is used to select which particle system number to apply the ParticleInstance modifier to, when the mesh which has the particle system on it has more than one of these. The PSYS field can have a value between **1** and **10**. It is possible to select any of the ten particle system numbers, however a check will **not** be made with the underlying particle emitting object specified previously in the Ob field. If you select a particle system number which does not exist on the particle emitting object, then the particles on the emitting mesh will keep their normal particle shapes â← Â no warning will be given that the chosen particle system does not exist on a particular particle emitting mesh.
As an example, below is a single plane mesh with two areas (the first area shown in red and the second in white), with different particle systems applied to each area. The left side using a ParticleInstance modifier which has the shape of a sphere and the right side having a ParticleInstance modifier which has the shape of a cube.



Render showing a single Plain mesh object assigned to two different vertex groups and each of those vertex groups is assigned a separate and independent particle system, with each particle system being assigned a different ParticleInstance modifier. In the case shown the ParticleInstance modifiers are a sphere and a cube.
Example Blend file

Normal

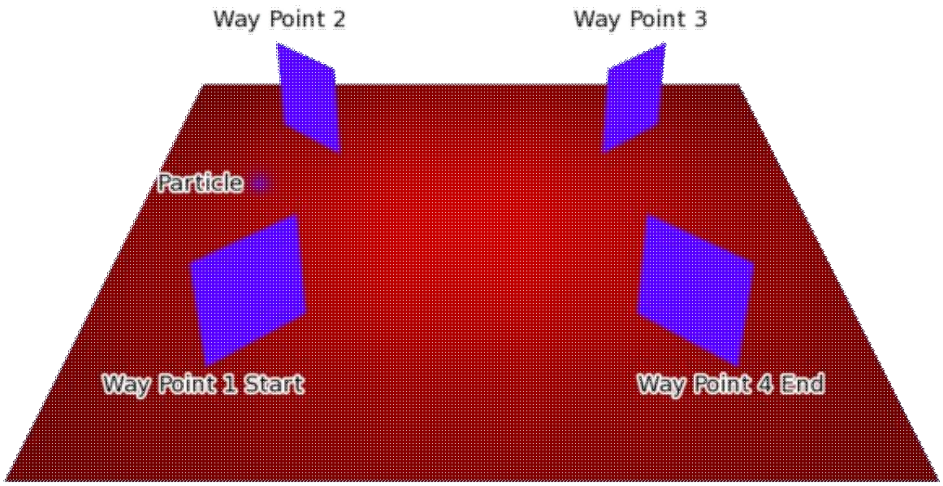
When selected, the Normal button tells the ParticleInstance modifier to draw instances of itself wherever normal particle types are emitted from the underlying particle system. So if the current ParticleInstance modifier is a sphere shape, when normal particles are emitted they will be spheres.

Children

When selected, the Children button tells the ParticleInstance modifier to draw instances of itself wherever children/child particles are emitted/used on the underlying particle system. So if the current ParticleInstance modifier is a sphere shape, when children/child particles are emitted they will be spheres.

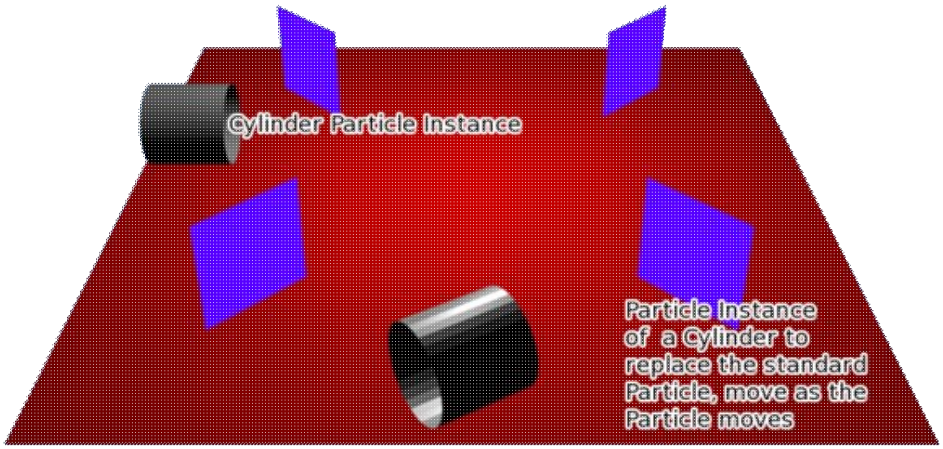
Path

The Path button tries to make the underlying mesh object of the ParticleInstance modifier deform its mesh shape in such a way as to try and match the path traveled by the particles/hair strands of the system associated with it.
For example, below is a screen shot showing the path of a single keyed particle as it travels its way through each of the different way points **1** to **4** (target particle systems), when it reaches way point **4** the particle dies and ends its journey.



Keyed particle following way points (showing one particle).
Example Blend file

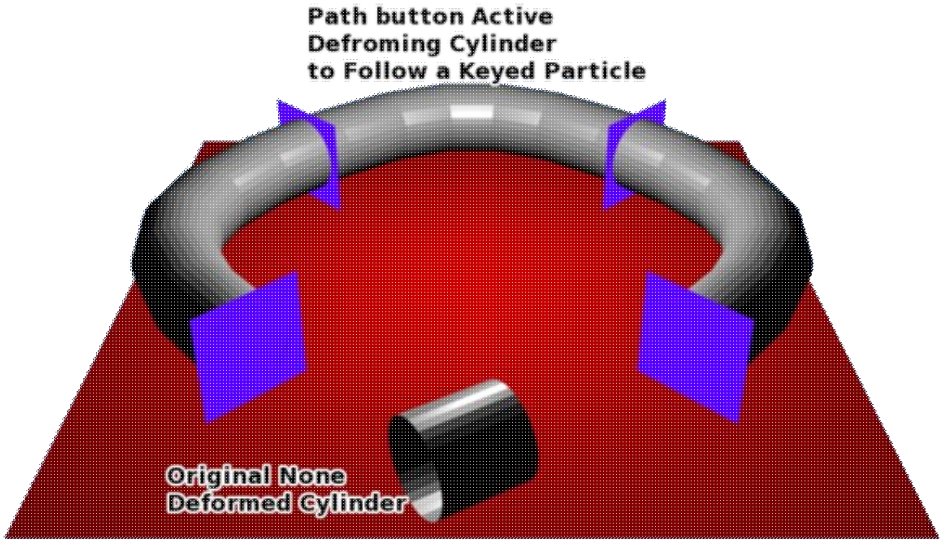
When a ParticleInstance modifier is added to a cylinder object and then associated with the way point particle system, the particle position is copied by the cylinder and placed at the particles position. So the mesh object follows the location of the particle. The cylinder does not alter any of its other properties when following the particle, only the cylinders location gets altered, shape and rotation do not get altered. See screenshot below:



Keyed particle following way points showing a mesh object (ParticleInstance modifier) in place of the original particle.

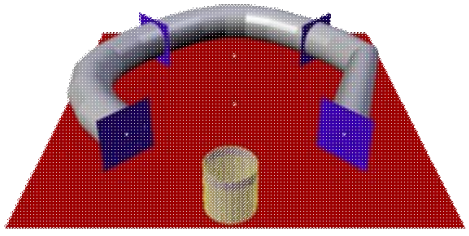
Example Blend file

Both of the above examples had the ParticleInstance modifier Path button deactivated.
When the Path button is activated the effect can be seen in the screenshot below:

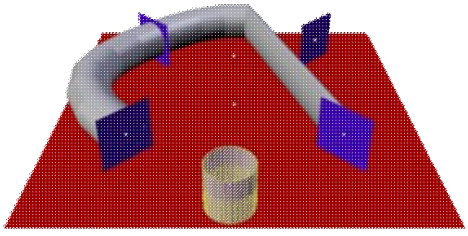


Keyed particle following way points showing a mesh object (ParticleInstance modifier) in place of the original particle, that is also being deformed to fit the travel path of the original particle.
Example Blend file

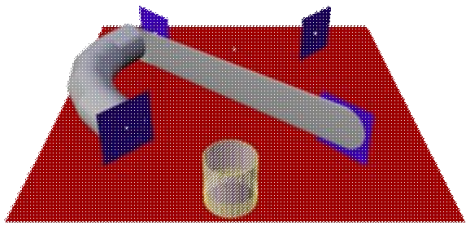
Instead of the cylinder location just following the position of the particle (and not altering its shape), the cylinder tries to fit its mesh to the shape of the path followed by the particle.
The mesh geometry of the object which is trying to deform can have an impact on how well the deformation is carried out. In the case of the cylinder, it has many loop cuts along its length so that it can bend at those points to deform along the particle path.
For example here is the same scene with the number of loop cuts along the length of the cylinder reduced, showing the effect on the deformation of the cylinder along the particle path.



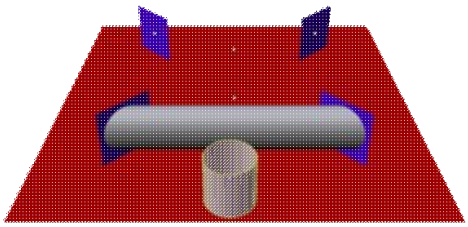
The cylinder has most of its edge loops so most of the path deform is very regular apart from at the very end of the curve.



The cylinder has some of its edge loops removed so the path of the deform starts to become less regular.

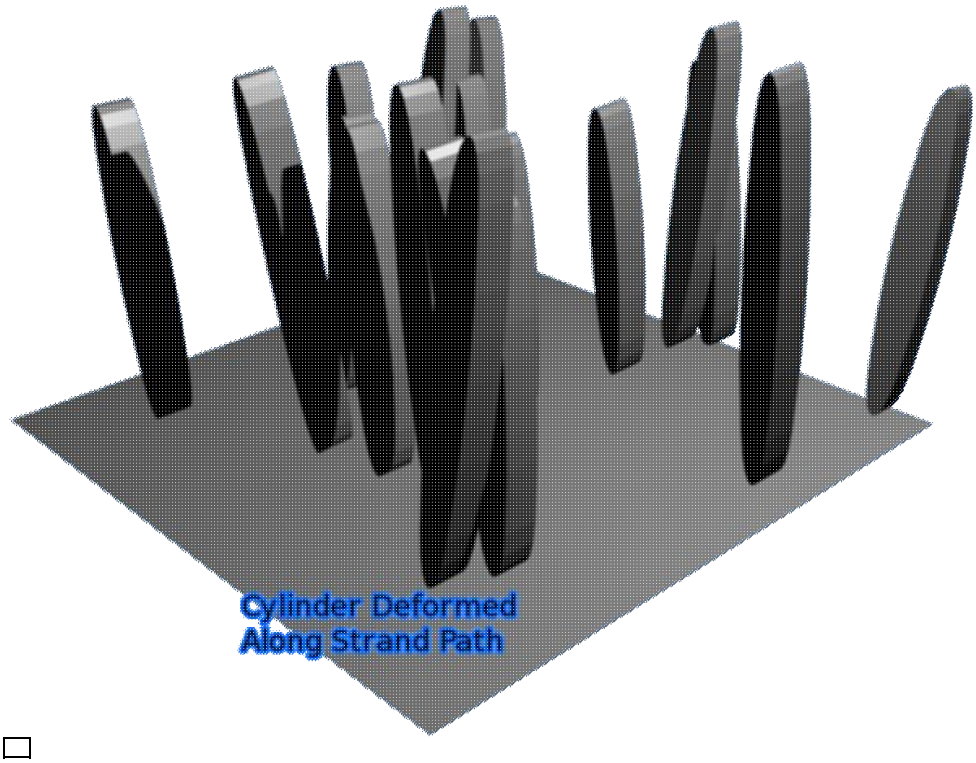


Now the deform path is very rough.



At this point there arenâ ← t any vertices to bend the cylinder to follow the path, and instead the cylinder just goes directly to the last way point **4**.

Once all the extra edge loops around cylinder are removed so that there is only the top and bottom vertices left, meaning that the cylinder doesnâ ← t have enough geometry to bend, in that case it cannot follow the path of the particle, so it just goes from the start way point **1** to the ending way point **4**.
The ParticleInstance modifier Path button works for hair (strand) particles as well as with keyed particles. In this case the mesh of the ParticleInstance modifier will follow the length and profile of the hair strands paths.
Below is a screenshot showing the effect of the Path button on hair:



Strand with a ParticleInstance modifier associated with it and deforming the cylinder along the hair profile.
Example Blend file

Strands when they are generated instantly die when created so for the Path button to be of any use, you must also have the Dead button activated. Otherwise the path a mesh took will not be visible!

Note

Thanks to Soylentgreen for explaining how the Path button works â← Â without his help, I would still have been completely lostâ←¦ -- Terrywallwork -- 6th Nov 2008.

Unborn

When selected, the Unborn button tells the ParticleInstance modifier to draw instances of itself wherever unborn particles will be emitted/used on the underlying particle system. So if the current ParticleInstance modifier is a sphere shape, when unborn particles are present they will be spheres.

Alive

When selected, the Alive button tells the ParticleInstance modifier to draw instances of itself wherever alive particles will be emitted/used on the underlying particle system. So if the current ParticleInstance modifier is a sphere shape, when alive particles are present they will be spheres.

Dead

When selected, the Dead button tells the ParticleInstance modifier to draw instances of itself wherever dead particles will occur on the underlying particle system. So if the current ParticleInstance modifier is a sphere shape, when dead particles are present they will be spheres.

See Also

- The whole particle section

Build Modifier

Mode: Object mode

Panel: Modifiers (Editing context, F9)

Description

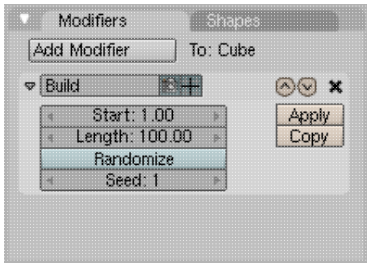
The Build modifier causes the faces of the mesh object to appear, one after the other, over time. If the material of the mesh is a halo rather than a standard one, then the vertices of the mesh, not the faces, appear one after another.

By default, faces (or vertices) appear in the order in which they are stored in memory (by default, they order of creation). These orders can be altered by selecting the object andâ←¦

- For vertices, in Edit mode, using the Xsort button (Mesh Tools panel) to reorder them along the X-axis (TODO: Once the new modelling chapter is published, add a link to Doc:Manual/Modeling/Meshes/Editing/Misc#Vertex List Ordering --Mont29 15:51, 5 February 2010 (UTC)).
- For faces, in Object mode, pressing CtrlF to sort them, based on different elements (TODO: Once the new modelling chapter is published, add a link to Doc:Manual/Modeling/Meshes/Editing/Misc#Face List Ordering --Mont29 15:51, 5 February 2010 (UTC)).

Never press the Apply button in this modifier panel if you want to keep the build animation!

Options



Build modifier panel.

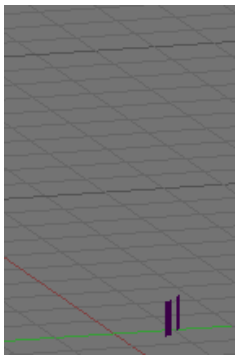
Start
: The start frame of the building process.

Length
: The number of frames over which to build up.

Randomize
: Randomizes the order that the faces are built up.

Seed
: The random seed. Change this to get a different â← randomâ← order of appearing â← Â this order being always the same for a given (seed, object) set.

Example



Small Build modifier example.
Imagine a city being built right before your eyes. How about a fountain of halos?

This example shows a simple tube, beauty short subdivided a few times, with the modifier turned on, from frames **1** to **100** (in steps of **10**), with no additional animation.

Bevel Modifier

Mode: Any mode

Panel: Modifiers (Editing context, F9)

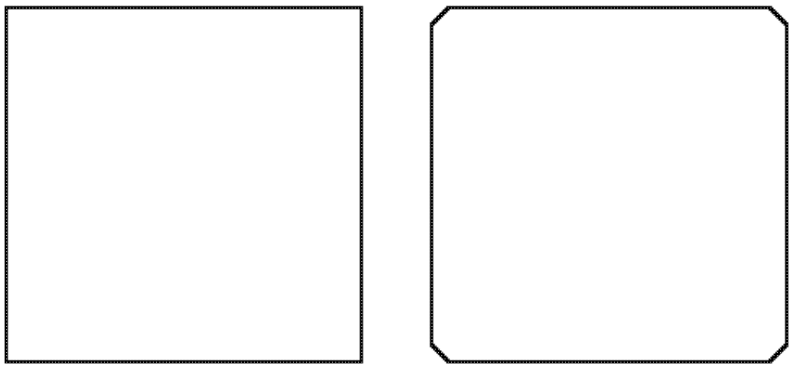
Description

The Bevel modifier adds the ability to bevel the edges of the mesh it is applied to, allowing control of how and where the bevel is applied to the mesh.

What is a Bevel?

The definition of a bevel, according to wikipedia.org, is:

> Â«Â A beveled edge refers to an edge of a structure that is not perpendicular (but instead often at 45 degrees) to the faces of the piece. The words bevel and chamfer overlap in usage; in general usage they are often interchanged, while in technical usage they may sometimes be differentiatedâ←¦Â Â»



Unbeveled square. Beveled square.

The picture (*Unbeveled square*) shows a square which has unbeveled edges as the angles between the corners of the square are 90Â° (perpendicular). The picture (*Beveled square*) shows a square which has beveled corners.



Default bevel.

Although the two pictures above show 2D squares, the Blender Bevel modifier can work on both 2D and 3D meshes of almost any shape, not just squares and cubesâ ←¦

The picture (*Default bevel*) shows a Blender 3D cube whith a bevel applied using just the default Bevel modifier settings.

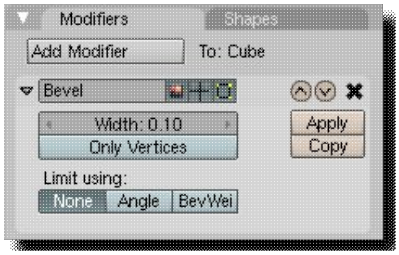
Setting Bevel Weight

When certain bevel options (such as BevWeight in the Bevel modifier) are set, the mesh will need to have bevel weights assigned to it, otherwise the bevel will not be applied.

Bevel weights can be applied to selected edges of a mesh by:

1. Switching to Edit mode.
2. Selecting the edge or edges of the mesh that you wish to assign a bevel weight to.
3. Pressing Ctrlâ § ShiftE or going to menu entry Mesh Â» Edges Â» Adjust Bevel Weight.
4. Then by moving the mouse around, or entering a value directly at the keyboard between **-1.0** to **+1.0**, you should be able to change the default bevel weight. You will be able to see the current value of the bevel weight change at the bottom status area of the 3D Viewport.
5. If the mesh already has a Bevel modifier with the BevWeight button enabled, as you alter the bevel weights, you should see the bevel amount changing in real time.

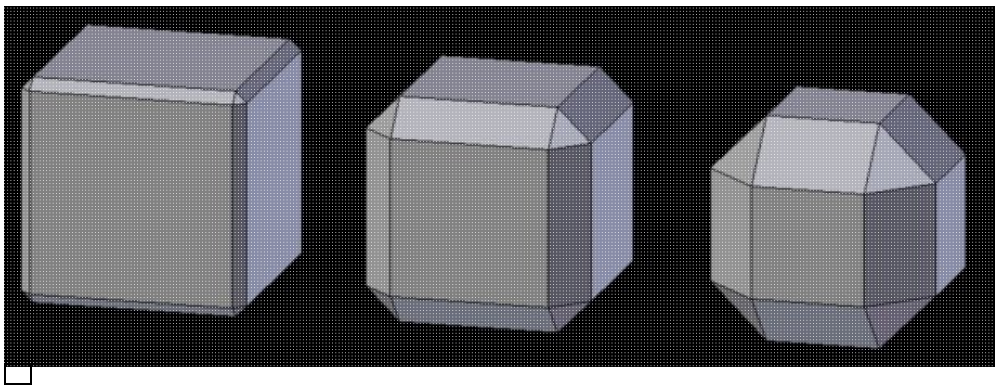
Options



Bevel modifier panel.
The Bevel modifier panel is a reasonably uncluttered panel and, for the most part, intuitive. That said, here is a description the settings contained within this panel:

Width

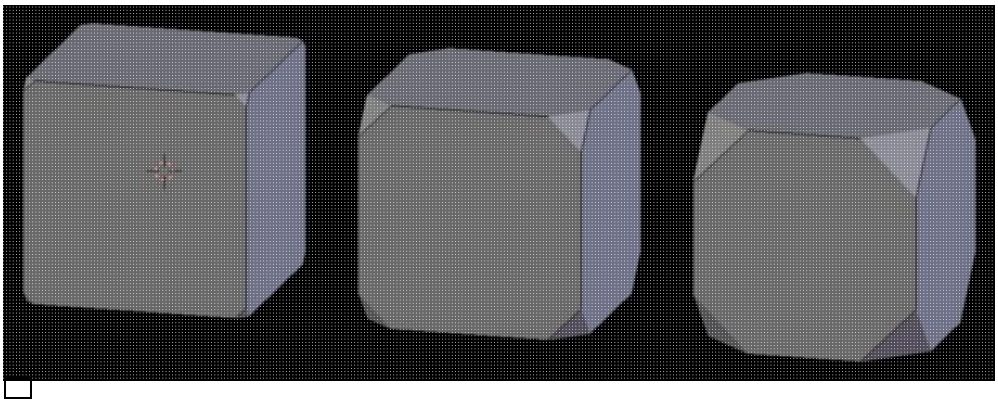
The Width numeric field controls the width/amount of the bevel applied to the base mesh. It can range from **0.0** (no bevel applied) to **0.5** (Blender Units). This value is in fact the â ← backing upâ ← of the two new edges relatively to the original (beveled) one, along the two concerned faces.



Three Cubes with **0.1**, **0.3** and **0.5** bevel Widths.

Only Vertices

The Only Vertices button alters the way in which a bevel is applied to the mesh. When it is active, only the areas near vertices are beveled, the edges are left unbeveled.



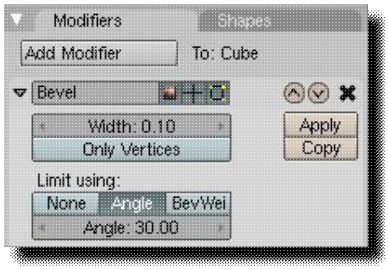
Three cubes with **0.1**, **0.3** and **0.5** bevel Widths, with Only Vertices option enabled.

Limit Using:

This section of the Bevel modifier is used to control where and when a bevel is applied to the underlying mesh. The first row of three buttons (mutually exclusive) controls the algorithm used, and might add some extra options.

None

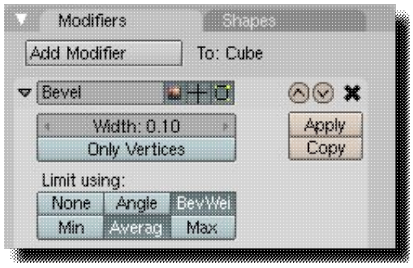
This button will apply the Bevel modifier to the whole underlying mesh, without any way to prevent the bevel on some edges/vertices.



Bevel modifier with the Angle limit enabled.

Angle

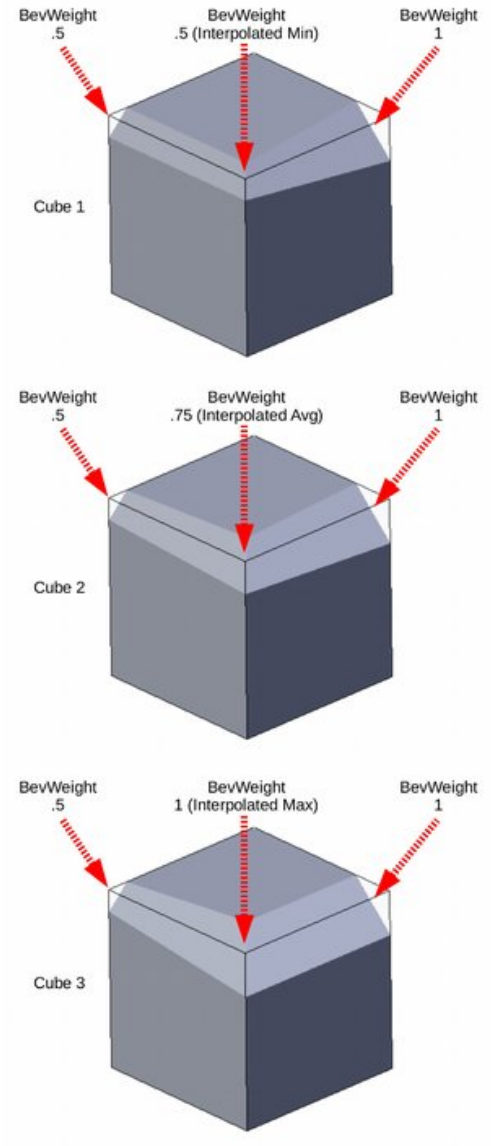
This button will only bevel edges where faces make sharp angles. When selected, it displays the Angle numeric field, used to set the angle above which an edge will be beveled (it is in fact the complementary angle, i.e. `180Â°-(angle between faces)`). When the angle between meeting faces is less than the angle in the slider box, the bevel on those specific edges will not be applied. Similarly, when the angle between two edges is less than this limit, the vertex is not beveled.



Bevel modifier with BevWeight button active.

BevWeight

This button will force the Bevel modifier to take into account any bevel weights that may or may not be assigned to the edges of the underlying mesh. When the BevWeight button is active, and at least some of the edges are bevel weighted, the Bevel modifier will bevel each edge according to its weight.
When the BevelWeight button is active, three extra buttons appear underneath it, named Min, Average and Max. They control how the bevel weighted influences are calculated and assigned at places where two or more bevel weighted edges meet, that have different weights. Obviously when there are several bevel weighted edges which are directly linked/connected to each other, there needs to be a way to determine which particular bevel weight gets assigned at the point of contact between them (interpolation vertex point in the diagrams of the three cubes, below).



Bevel modifier with BevWeight limit, using Min, Average and Max interpolation modes.

Min
: When several bevel weighted edges with different weights meet, the interpolation/contact point between them will take the smaller of the bevel weights as the one for this point (see `Cube 1` to the right).

Average
: When several bevel weighted edges meet, the interpolation/contact point between them will take the average value of the bevel weights as the one for this point (see `Cube 2` to the right).

Max
: When several bevel weighted edges meet, the interpolation/contact point between them will use the larger of the bevel weights as the one for this point (see `Cube 3` to the right).

Note
Thanks to Xalt for making clear how the limit using BevWeight Min, Average and Max buttons affect underlying meshes. -- Terrywallwork -- 28th June 2008.

See Also

- The old and deprecated Bevel tool.

Boolean Tools

Mode: Object mode (meshes only)

Panel: Modifiers (Editing context, F9)

Hotkey: W

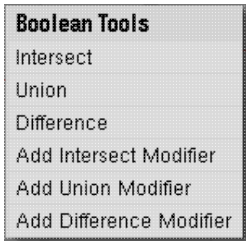
Menu: Object Â» Boolean Operation...

Description

Boolean operations are a method of combining or subtracting solid objects from each other to create a new form, a new object. Boolean operations in Blender only work on two Mesh type objects, preferably ones that are solid, or closed, with a well defined interior and exterior surface. If more than two mesh objects are selected only the *active* and previously selected object are used as operands. The boolean operations also take materials and UV-textures into account, producing objects with material indices or multi UV-mapped objects.

Options

Using the Boolean menu (W in Object mode) presents the following options:



Boolean operations.

Intersect
: Creates a new object whose surface encloses the volume *common to both* original objects.

Union
: Creates a new object whose surface encloses the *total volume of both* original objects.

Difference
: The only operation in which the order of selection is important, the active object is subtracted from the selected object. That is, the resulting object surface encloses a volume which is the volume belonging to the selected *and inactive* object, but *not* to the *active* one.

Add Intersect Modifier
: A shortcut that applies a Boolean modifier and selects Intersect in one step.

Add Union Modifier
: A shortcut that applies a Boolean modifier and selects Union in one step.

Add Difference Modifier

Guess what? A shortcut that applies a Boolean modifier and selects Difference in one step.

Boolean Modifier

Mode: Any mode

Panel: Modifiers (Editing context, F9)

Description

This sub-panel appears in the Editing context panel group which is accessed using F9 or clicking  button in the Buttons window. It is part of the Modifiers parent panel. For further information about the common panel components see the Interface section on modifiers.

The downside of using the direct boolean commands is that in order to change the intersection, or even apply a different operation, you need to remove the new object and redo the command. In contrast, one can use a Boolean modifier for great flexibility and *non-destructive editing*. As with all modifiers, the booleans can be enabled/disabled or even rearranged in the stack. In addition, you can move the operands and see the boolean operation applied interactively in real time!

Caution
If the objectsâ ← meshes are too complex you may be waiting a while as the system catches up with all the mouse movements. Turning off display in the 3D View in the modifier panel can improve performance.

To get the final, â ← definitiveâ ← object from this modifier (like with the direct boolean tools) you need to â ← applyâ ← the modifier using the modifierâ ← s Apply button, and to see the result you need to move the remaining operand away or switch to local view / NumPad. Until you apply the modifier, the objectâ ← s mesh will not be modified. When you apply the boolean modifier you are notified that any mesh sticky information, animation keys and vertex information will be deleted.

Warning
There is an important difference between using boolean tools and applying a Boolean modifier: the first one **creates a new object**, whereas the second **modifies its underlying objectâ ← s mesh**. This means that when you apply a boolean modifier, **you â ← loseâ ← one of your operands**! Make a duplicate (â § ShiftD) of the object bearing the modifier, if you want to keep an â ← originalâ ← version.

However, you may not want to apply your Boolean modifier: it is indeed a very nice tool to perform some animated effects/deformations, as it is re-evaluated for each frame. This allow you e.g. to have a sort of â ← force fieldâ ← deforming around an object when this one penetrates it (see example below). Note that you can set objects to be invisible on render, which might be very useful with one of the operands of a boolean operation!

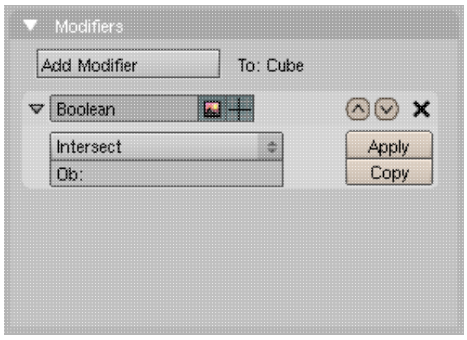
From Blender 2.49, Boolean modifiers need no more to be at the top of the modifiers stack. This was quite a heavy limitation, as it meant that you couldnâ ← t use more than one Boolean modifier at a time on a same mesh, and that you could only work on â ← rawâ ← mesh data (i.e. couldnâ ← t use e.g. the result of a Subsurf modifier as a boolean operand â ← unless you applied it first, of course).

Now, you can place a Boolean modifier where you want in the stack. You can also have multiple Boolean modifiers on the same object. And last but not least, the â ← second operandâ ← object (the one you enter the name in the Ob field of the modifier) will also be used â ← all-modifiers-appliedâ ← (e.g. you can use its subsurfed form).

Warning

All these new possibilities are great. But be careful: the boolean operations are still a complex process, which may take a very long time to complete on operands with many vertices. So take attention to all modifiers placed *before* the Boolean one, especially to Subsurf ones!

Options



Modifiers panel with Boolean modifier activated.

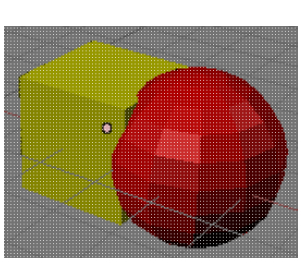
Operation
: This drop-down lists available boolean operation types (Intersect, Union or Difference).

Ob
: The name of the object to be used as the second operand to this modifier.

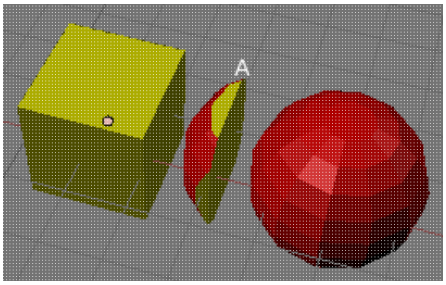
Examples

Intersect

The cube and the sphere have been moved to reveal the newly created object (â← Aâ←). Each face of the new object has the material properties of the corresponding surface that contributed to the new volume based on the Intersect operation.



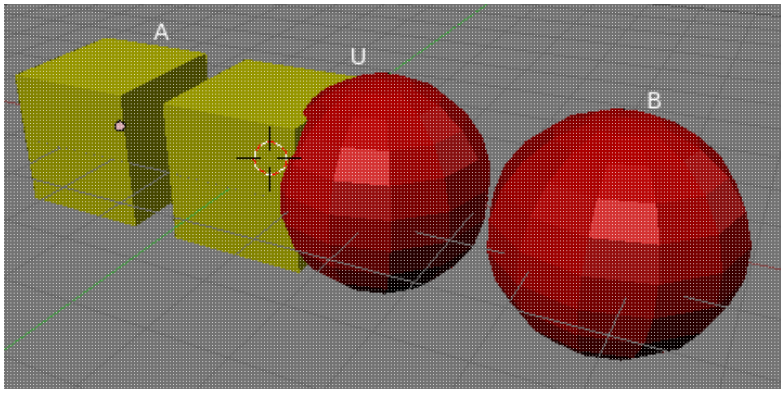
Before Intersect.



After Intersect.

Union

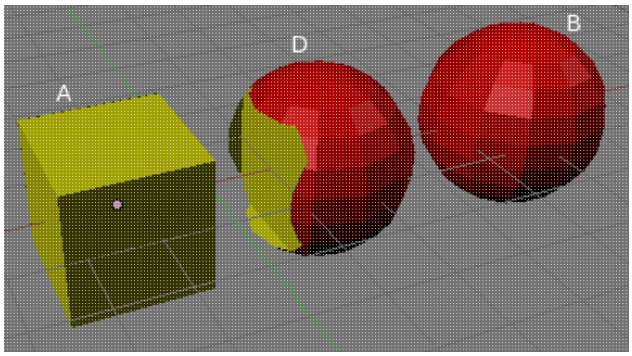
The cube (â← Aâ←) and the sphere (â← Bâ←) have been moved to reveal the newly created object (â← Uâ←). â← Uâ← is now a single mesh object and the faces of the new object have the material properties of the corresponding surface that contributed to the new volume based on the Union operation.



Union example.

Difference

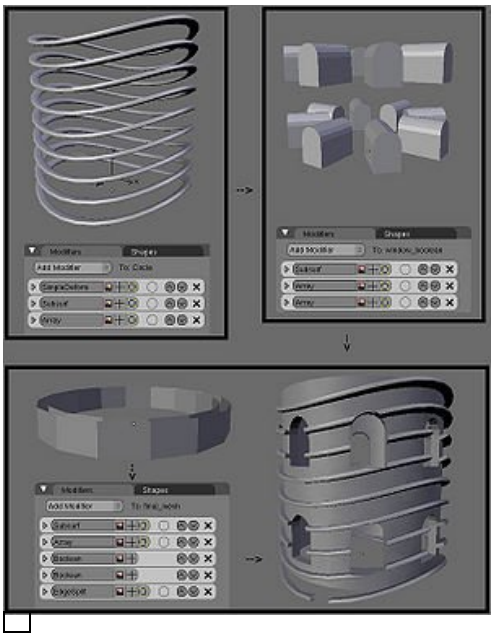
The Difference of two objects is (as the â← standardâ← subtraction) not commutative in that the *inactive* object minus the *active* object (behavior of this operation) does not produce the same as *active* minus *inactive*. The *active* cube (â← Aâ←) has been subtracted from an *inactive* sphere (â← Bâ←), and both have been moved to reveal the newly created object (â← Dâ←). â← Dâ← is now a single mesh object and the faces of the new object have the material properties of the corresponding surface that contributed to the new volume based on the Difference operation. Dâ← s volume is less than Bâ← s volume because it was decreased by subtracting part of the cubeâ← s volume.



Difference example.

More complex examples with Boolean modifiers

Multiple Modifiers



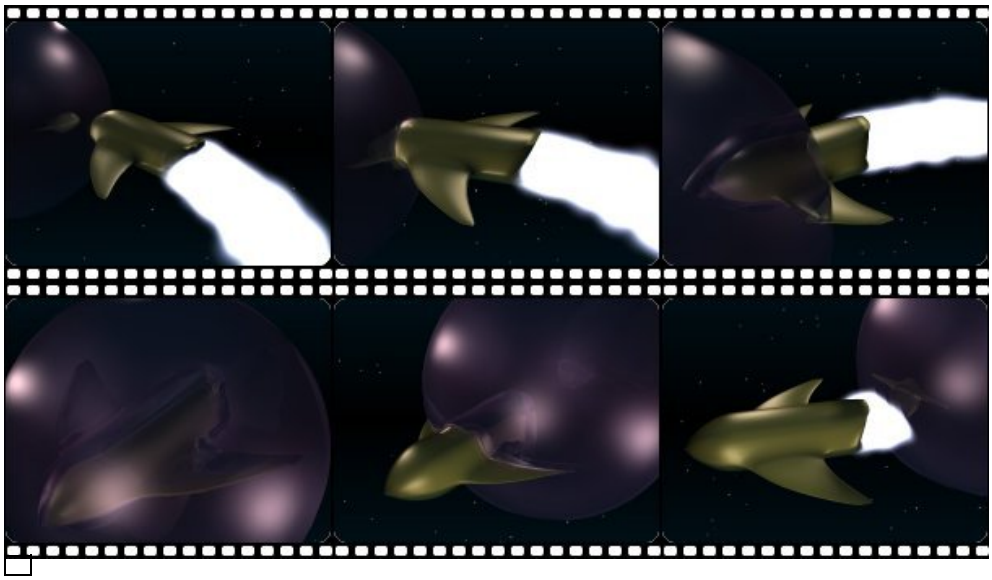

A complex object modeled using multiple Boolean modifiers.

The process to model the complex object shown to the right.
(*The process to model the complex object shown to the right*) illustrates how to model quite a complex mesh object, using only three basic forms, and several modifiers. As usual, click on the picture to enlarge.

(*A complex object modeled using multiple Boolean modifiers*) is the result of this modeling process. You can download the demo blender file here.

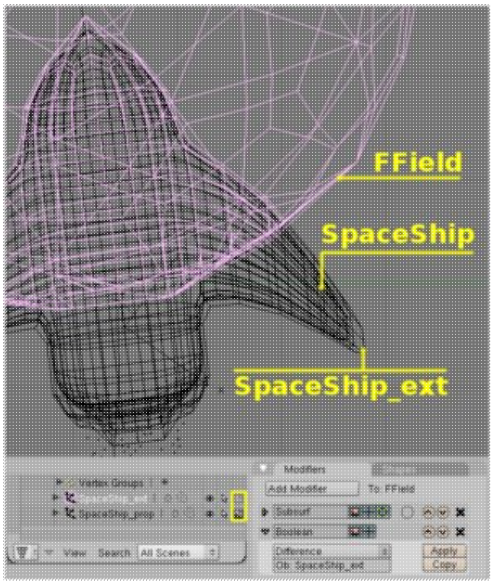
Here is a video showing the new Blender 2.49 workflow:

Animation



An animated boolean deformation.

(*An animated boolean deformation*) shows some key-frames of the boolean-deformation of a â← force fieldâ← sphere penetrated by a sort of space swallow. It tries to demonstrate the possibilities of Boolean modifiers in animations.



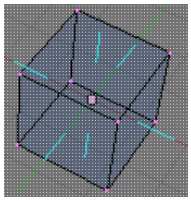
The schema of the effect.

(*The schema of the effect*) summarizes the key points of the boolean-deformation effect. There is one small trick here: one of the boolean operands (â← `SpaceShip_ext`â←) is not rendered (the third toggle control, at the right of the Outliner window, in Outliner display), which allows us to keep the â← `FField`â← surface a bit away from the real â← `SpaceShip`â← (else, we would have a lot of artifacts!).
â← `SpaceShip_ext`â← is a linked-duplicate (AltD) of the â← `SpaceShip`â← object, just slightly up-scaled and parented to the â← originalâ← .

You can download the blender file here.

And here is the complete video of this anim (Media:manual-modifiers_boolean_exvideo.avi):

Technical Details



Visible normals.
The boolean operations rely heavily on the surface normals of each object and so it is very important that the normals are defined properly and consistently. This means each objectâ← s normals should point outward. A good way to see the objectâ← s normals is to turn on the visibility of normals using the Mesh Tool More panel (Buttons window, Editing context â← Â F9) and clicking Draw normals. The normals are only visible while in Edit mode. (*Visible normals*) is an example of a cube with its normals visible.

In the case of *open* objects, that is objects with holes in the surface, the interior is defined mathematically by extending the boundary faces of the object to infinity. As such, you may find that you get unexpected results for these objects. A boolean operation never affects the original objects, the result is always a new object.

Warning
This is **not true** with the Boolean modifiers: when they are applied, they modify their owner object, and do not create a new one!

Some operations will require you to move the operands or switch to local view (/ NumPad) to see the results of the boolean operation.

Limitations & Workarounds

The number of polygons generated can be very large compared to the original meshes, especially when using complex concave objects. Furthermore, the polygons that are generated can be of poor quality, for example, very long and thin and sometimes very small. Try using the Decimate modifier (Editing context, F9) to fix this problem.

Sometimes the boolean operation can fail with a message saying â← `An internal error occurred -- sorry`â← . If this occurs, try to move or rotate the objects just a very small amount and try again.

Decimate Modifier

Mode: Object mode

Panel: Modifiers (Editing context, F9)

This sub-panel appears in the Editing context panel group which is accessed using F9 or clicking button in the Buttons window. This sub-panel is part of the Modifiers parent panel. For further information about the common panel components see the modifier stack section.

Description

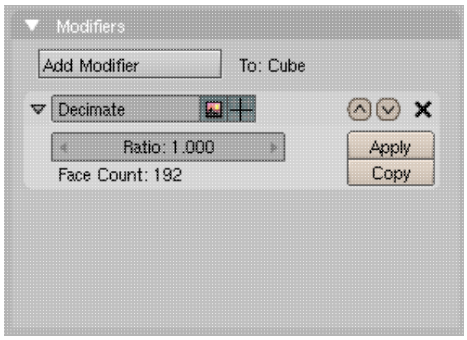
The Decimate modifier allows you to reduce the vertex/face count of a mesh with minimal shape changes. This is not applicable to meshes which have been created by modeling carefully and economically, where all vertices and faces are necessary to correctly define the shape, but if the mesh is the result of complex modeling, with proportional editing, successive refinements, possibly some conversions from SubSurfed to non-SubSurfed meshes, you might very well end up with meshes where lots of vertices are not really necessary.

The Decimate modifier is a quick and easy way of reducing the polygon count of a mesh non-destructively. This modifier demonstrates of the advantages of a mesh modifier system because it shows how an operation, which is normally permanent and destroys original mesh data, can be done interactively and safely using a modifier.

Unlike the majority of existing modifiers, the Decimate modifier does not allow you to visualize your changes in Edit mode.

The Decimate tool only handles triangles, so each quadrilateral face is implicitly split into two triangles for decimation.

Options



Decimate modifier.

Ratio
: The ratio of faces to keep after decimation, from **0.0** (0%, all faces have been completely removed) to **1.0** (100%, mesh is completely intact, except quads have been triangulated).
As the percentage drops from **1.0** to **0.0**, the mesh becomes more and more decimated until it no longer visually looks like the original mesh.

Face Count (display only)

This field shows the number of remaining faces as a result of applying the Decimate modifier.

Examples

Simple plane

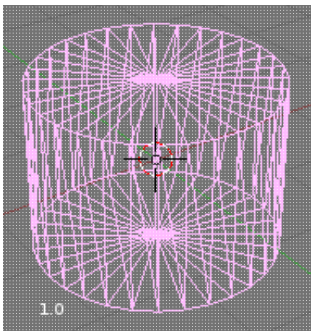
A simple example is a plane, and a 4x4 undeformed Grid object. Both render exactly the same, but the plane has one face and four vertices, while the grid has nine faces and sixteen vertices, hence lots of unneeded vertices and faces. The Decimate modifier allows you to eliminate these unneeded faces.

Decimated cylinder

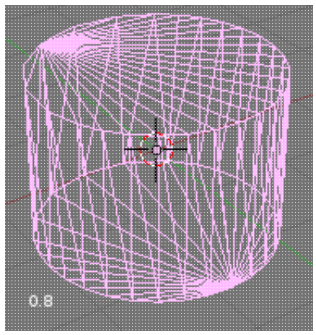
We take a simple example of decimating a cylinder using the default of **32** segments. This will generate a cylinder with **96** faces. When the Decimate modifier is applied, the face count goes up! This is because the modifier converts all quadrangles (*quads*) into triangles (*tris*) which always increases the face count. Each quad decomposes into two triangles.

The main purpose of the Decimate modifier is to reduce mesh resources through a reduction of vertices and faces, but at the same time maintain the original shape of the object.

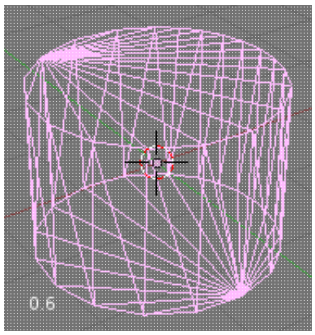
In the following pictures, the percentage dropped in each successive image, from **100%** to **5%** (a ratio of **0.05**). Notice that the face count has gone from **128** to **88** at a ratio of **0.6** (**60%**) and yet the cylinder continues to look very much like a cylinder and we discarded **40** unneeded faces.



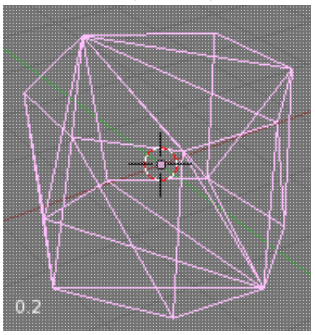
Ratio: **1.0** (**100%**). Faces: **128**.



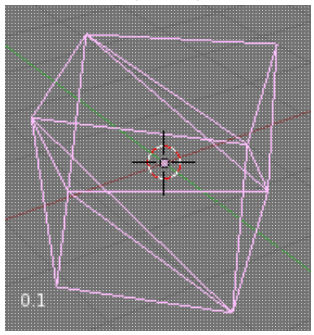
Ratio: **0.8** (**80%**). Faces: **102**.



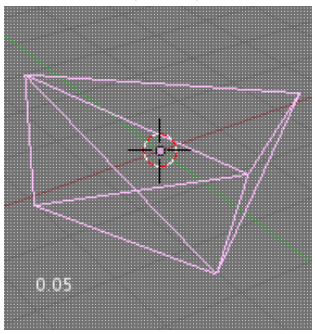
Ratio: **0.6** (**60%**). Faces: **88**.



Ratio: **0.2** (**20%**). Faces: **24**.



Ratio: **0.1** (**10%**). Faces: **12**.



Ratio: **0.05** (**5%**). Faces: **6**.

As you can see when the ratio has reached **0.1**, the cylinder looks more like a cube. And when it has reached **0.05** it doesnâ ← t even look like a cube!

Once you have reached the face count and appearance you were looking for you can Apply the modifier. If you want to convert as many of the tris back to quads to reduce mesh resources further you can switch to Edit mode, select all vertices (A), and hit AltJ.

High resolution landscape



Decimated landscape, top: original; middle: lightly decimated; bottom: heavily decimated.
Decimated landscape, top: original; middle: lightly decimated; bottom: heavily decimated. shows a landscape generated via a careful application of the Noise tool (or, nowadays, of the Displace modifierâ←¦), on a quite vast grid. On top, the result for the original mesh and below, two different levels of decimation. To the eye the difference is indeed almost unnoticeable, but as the vertex count goes down there is a huge gain.

Displace Modifier

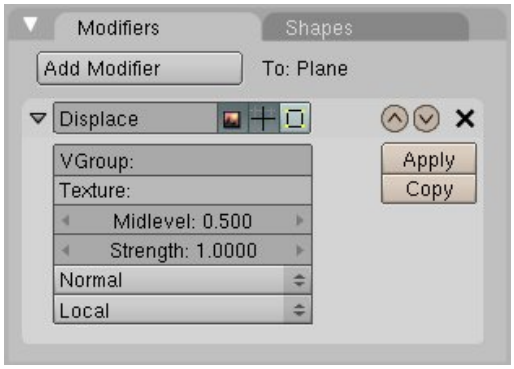
Mode: Any mode

Panel: Modifiers (Editing context, F9)

Description

The Displace modifier displaces vertices in a mesh based on the intensity of a texture. Either procedural or image textures can be used. The displacement can be along a particular local axis, along the vertex normal, or the separate RGB components of the texture can be used to displace vertices in the local X, Y and Z directions simultaneously.

Options



Displace modifier.

VGroup
: The name of a vertex group which is used to control the influence of the modifier.
If VGroup is empty, the modifier affects all vertices equally.

Texture
: The name of the texture datablock from which the displacement for each vertex is derived.
If this field is empty, the modifier will be disabled.

Midlevel
: The texture value which will be treated as no displacement by the modifier. Texture values below this value will result in negative displacement along the selected direction, while texture values above this value will result in positive displacement. This is achieved by the equation `(displacement)Â =Â (texture value)Â -Â Midlevel`.
Recall that color/luminosity values are typically between **0.0** and **1.0** in Blender, and not between **0** and **255**.

Strength
: The strength of the displacement. After offsetting by the Midlevel value, the displacement will be multiplied by the Strength value to give the final vertex offset. This is achieved by the equation `(vertex_offset)Â =Â (displacement)Â Ã Â Strength`.
A negative strength can be used to invert the effect of the modifier.

Direction

The direction along which to displace the vertices.
Can be one of the following:

- ◊ X â← displace along local X axis.
- ◊ Y â← displace along local Y axis.
- ◊ Z â← displace along local Z axis.
- ◊ RGB -> XYZ â← displace along local XYZ axes individually using the RGB components of the texture.
- ◊ Normal â← displace along vertex normal.

Texture Coordinates

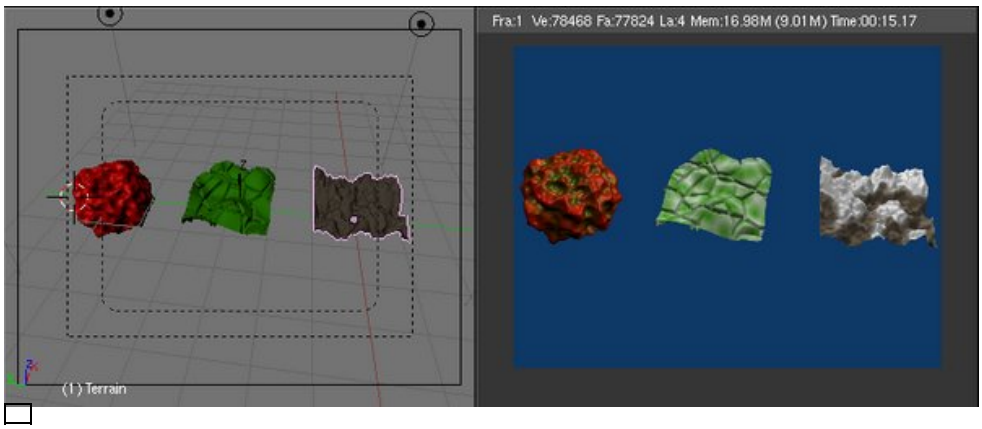
The texture coordinate system to use when retrieving values from the texture for each vertex.
Can be one of the following:

- ◊ UV â← take texture coordinates from face UV coordinates.
 UV Layer
 The UV coordinate layer from which to take texture coordinates.
 If the object has no UV coordinates, it uses the Local coordinate system. If this field is blank, but there is an UV coordinate layer available (e.g. just after adding the first UV layer to the mesh), it will be overwritten with the currently active UV layer.

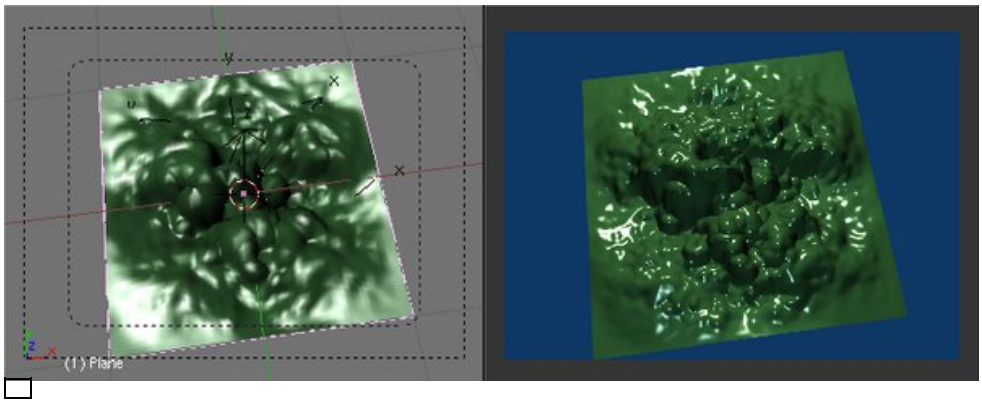
Note
Since UV coordinates are specified per face, the UV texture coordinate system currently determines the UV coordinate for each vertex from the first face encountered which uses that vertex; any other faces using that vertex are ignored. This may lead to artifacts if the mesh has non-contiguous UV coordinates.

- ◊ Object â← take the texture coordinates from another objectâ← s coordinate system (specified by the Ob field).
 Ob
 The object from which to take texture coordinates. Moving the object will therefore alter the coordinates of the texture mapping.
 If this field is blank, the Local coordinate system is used.
- ◊ Global â← take the texture coordinates from the global coordinate system.
- ◊ Local â← take the texture coordinates from the objectâ← s local coordinate system.

Examples



Three different objects created with the Displace modifier.
Sample blend file



A slime animation created with the Displace modifier.
Sample blend file

See also

- Blender artists post: Mesh modifier: displace

EdgeSplit Modifier

Mode: Any mode

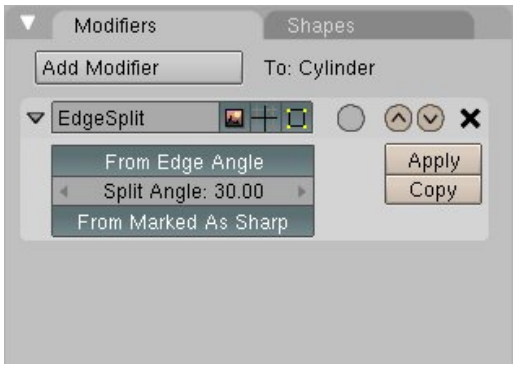
Panel: Modifiers (Editing context, F9)

Description

The EdgeSplit modifier splits edges within a mesh. The edges to split can be determined from the edge angle (i.e. angle between faces forming this edge), and/or edges marked as sharp.

Splitting an edge affects vertex normal generation at that edge, making the edge appear sharp. Hence, this modifier can be used to achieve the same effect as the Autosmooth button, making edges appear sharp when their angle is above a certain threshold. It can also be used for manual control of the smoothing process, where the user defines which edges should appear smooth or sharp (see Mesh Smoothing for other ways to do this). If desired, both modes can be active at once. The output of the EdgeSplit modifier is available to export scripts, making it quite useful for creators of game content.

Options



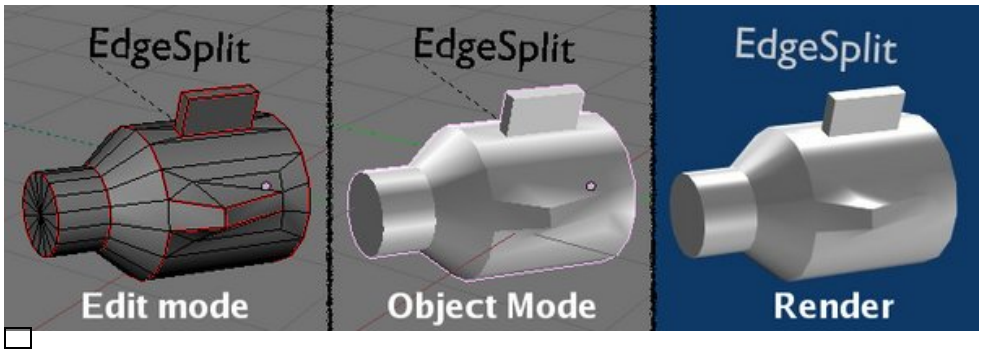
EdgeSplit modifier.

From Edge Angle
: If this button is enabled, edges will be split if their edge angle is greater than the Split Angle setting.
 - ◊ The edge angle is the angle between the two faces which use that edge.
 - ◊ If more than two faces use an edge, it is always split.
 - ◊ If fewer than two faces use an edge, it is never split.

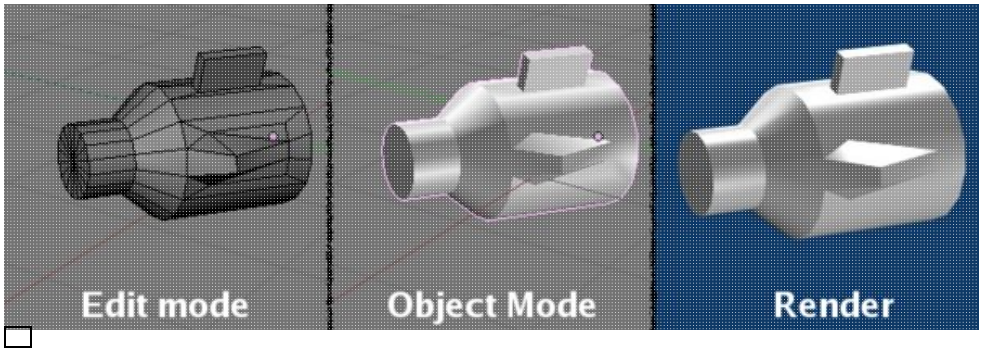
Split Angle
: This is the angle above which edges will be split if the From Edge Angle button is selected, from **0Â°** (all edges are split) to **180Â°** (no edges are split).

From Marked As Sharp
: If this button is enabled, edges will be split if they are marked as sharp, using the Edge Specials Â» Mark Sharp menu item (CtrlE, in Edit mode).

Examples



EdgeSplit modifier output with From Marked As Sharp selected.
Sample blend file



EdgeSplit modifier output with From Edge Angle selected.
Sample blend file

Explode Modifier

Mode: Any mode

Panel: Modifiers (Editing context, F9)

Description

The Explode modifier is used to alter the mesh geometry (by moving/rotating its faces) in a way that (roughly) tracks underlying emitted particles that makes it look as if the mesh is being exploded (broken apart and pushed outward).

For the Explode modifier to have a visible effect on the underlying mesh it has to be applied to a mesh which has a particle system on it, in other words it has to be a mesh which outputs particles. This is because the particle system on the mesh is what controls how a mesh will be exploded, and therefore without the particle system the mesh wont appear to alter. Also both the number of emitted particles and number of faces determine how granular the Explode Modifier will be. With default settings the more faces and particles the more detailed the mesh exploding will be, because there are more faces and particles to affect detachment/movement of faces.

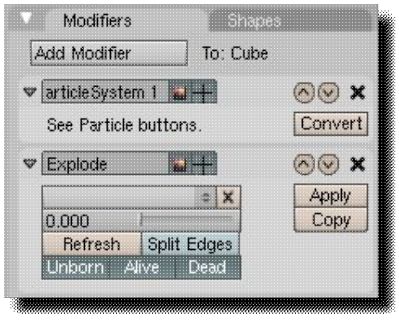
Here is a link to an Ogg Theora video showing a cube with a particle system and Explode modifier applied:

Media:Manual - Explode Modifier - Exploding Cube.ogg

Here is a link to the original Blender file which has an Exploding cube setup, just free the particle cache by pressing the Free Bake button in the Bake panel and then press the Animate button to see the animation:

Media:Manual - Explode Modifier - Exploding Cube.blend

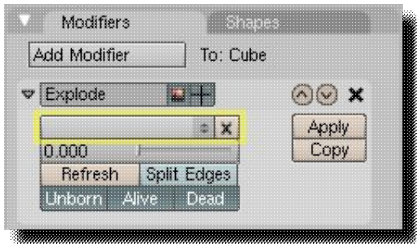
Options



Explode modifier panel with ParticleSystem modifier above it.

Stacking Order Importance

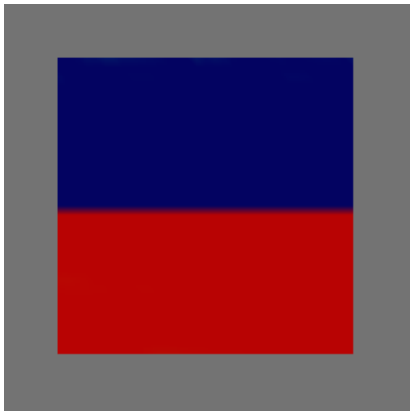
This modifier is highly affected by its position within the modifier stacking order. If it is applied before a ParticleSystem modifier it will not be affected by particles and therefore appear to do nothing. The ParticleSystem modifier must appear before the Explode modifier, because the first one has the information needed to drive the last one.



Explode modifier panel with â← Protect this vertex groupâ← field highlighted in yellow.

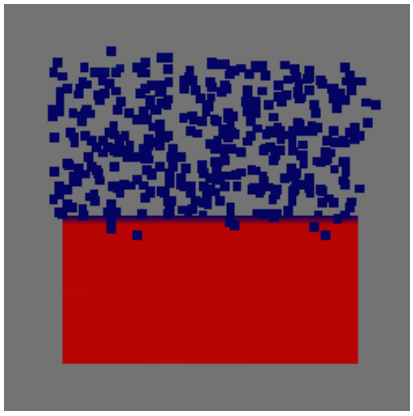
Protect this vertex group

If a mesh that has an Explode modifier on it also has vertex groups assigned to it, then this field will allow the selection of one of those vertex groups. This will force the Explode modifier to take into account the weight values assigned to the vertices of this group: it then must protect the faces from explosion, depending on the weights assigned their vertices (full protection for a weight of **1.0**, no protection for a weight of **0.0**).
It is also possible to have in between face protection: just assign them intermediate weights, between **0.0** and **1.0**.



Weight painted plane before Explode modifier is run.

To the right is a plane which has been weight painted in two different weight paints, red which represents a weight of **1.0**, and blue which represents a weight of **0.0**. The reason that weight painting is used is because whenever you weight paint something it automatically makes a new vertex group (if the plane doesnâ← t already have one), which can then be used with the â← Protect this vertex groupâ← field.



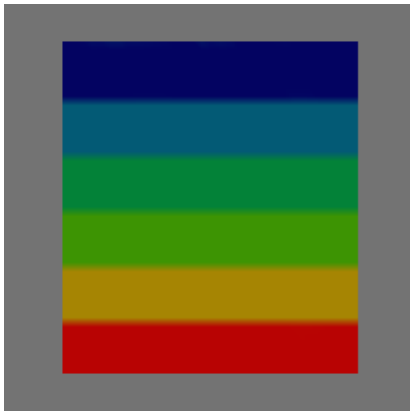
Weight painted plane after Explode modifier is run. Showing the red area unmodified and the blue area exploded.

To the right is the same plane while the Explode modifier is in the process of exploding the plane. As can be seen the weight painting means that the red (weight of **1.0**) faces are left unmodified, and the blue (weight of **0.0**) faces are modified.
Remember that by default weight painting makes a vertex group, that is what needs to be selected in the â← Protect this vertex groupâ← field for the weight painting to be taken notice of.
Here is a link to an Ogg Theora video showing a weight painted plane, painted with minimum and maximum weights to be animated with the Explode modifier:

Media:Manual_-_Explode_Modifier_-_Dual_Weighted_Plane.ogg

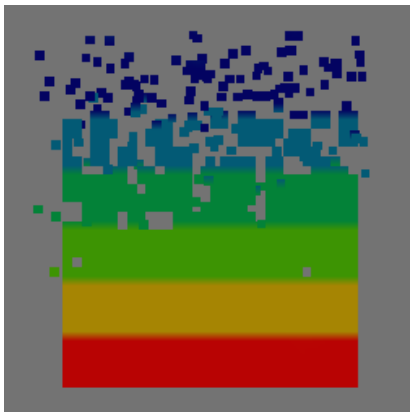
And here is a link to the original Blender file which the above movie was made from, just free the particle cache by pressing the Free Bake button in the Bake panel and then press the Animate button to see the animation:

Media:Manual_-_Explode_Modifier_-_Dual_Weighted_Plane.blend



Blended weight painted plane before the Explode modifier is run.

To the right is a plane which has got a blended weight paint on it, with red being a weight of **1.0** through to other lower weighted colors all the way to blue which is a weight of **0.0**.



Blended weight painted plane after the Explode modifier is run.

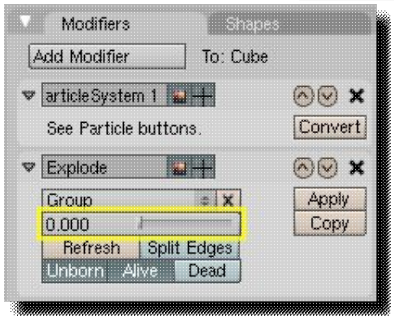
And now is the same blended weight painted plane after it has had the Explode modifier running on it. As can be seen the top part of the blended plane is highly affected by the Explode modifier, whereas, lower down the plane, where the weight painting gets higher in weight, is affected less and less by the

Explode modifier.
Here is a link to an Ogg Theora movie showing a blended weight painted plane, painted with multiple weights animated with the Explode modifier:

Media:Explode_Modifier_-_Graded_Weighted_Plane_-_Exploded.ogg

And here is a link to the original Blender file which the above movie was made from, just free the particle cache by pressing the Free Bake button in the Bake panel and then press the Animate button to see the animation:

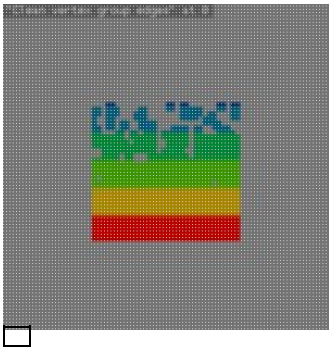
Media:Explode_Modifier_-_Graded_Weighted_Plane.blend



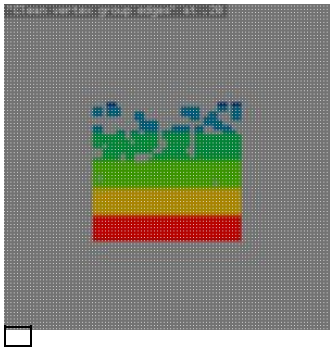
The â← Clean vertex group edgesâ← numeric slider field highlighted in yellow.

Clean vertex group edges

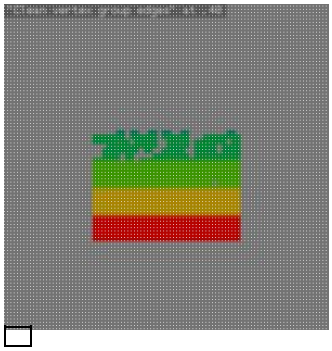
This numeric slider field appears to erode the exploded edges around a vertex groups edges (making edges less jaggied). The slider can have a value between **0.0** (no eroding of edges) and **1.0** maximum erosion of edges.
Below are some screenshots of a mesh that has had an Explode modifier applied to it with various values for â← Clean vertex group edgesâ← :



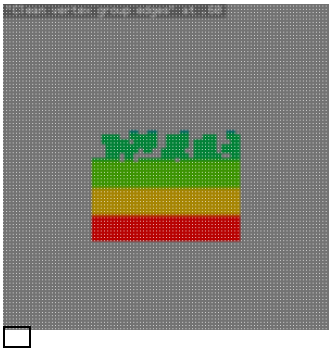
â← Clean vertex group edgesâ← to **0.0**.



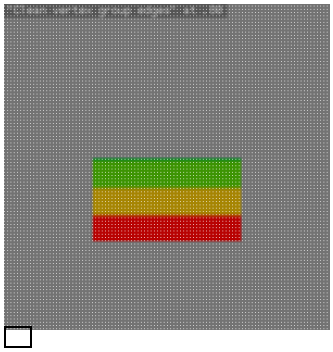
â← Clean vertex group edgesâ← to **0.2**.



â← Clean vertex group edgesâ← to **0.4**.



â← Clean vertex group edgesâ← to **0.6**.



â← Clean vertex group edgesâ← to **0.8**.

Caution

I have been unable to find out a definitive explanation of exactly how the â← Clean vertex group edgesâ← works on a mesh and what the slider value represents, so I have just used my best guess from looking at the result of using it and setting it with different values. If you have a more concrete definition of what is does please alter this text to reflect it or contact me and I will update the text. --Terrywallwork

Refresh (Recalculate faces assigned to particles)

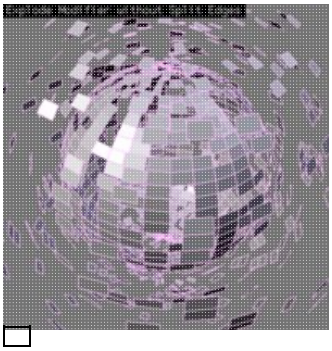
When certain changes are made to the underlying Explode modifier mesh (such as changing the vertex group weightings of faces), the faces displayed and the particles that influence certain faces can get out of sync. If this happens pressing the Refresh button will tell the Explode modifier to update all its calculations to take into account new settings.

As an example, if you have a mesh with faces that get modified by the end of an Explode modifier run, and you then alter those faces vertex group weights so they have a value of **1.0** (meaning they can no more get altered in future â← Â when the â← Protect this vertex groupâ← option is active), the Explode modifier will still show the old altered state of the faces until you click the Refresh button.

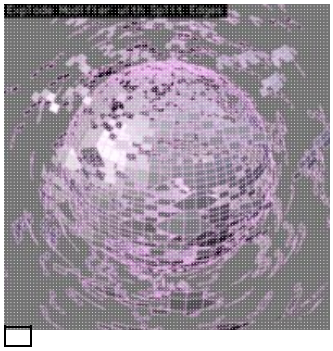
Split Edges

When this button is selected, faces which are exploded are split making them smaller and more fragmented, giving more faces for the Explode modifier to work with and usually giving better explode results.

Below are two screenshots of an Explode modifier in progress, one with no Split Edges option active and one with Split Edges activated:



Uvsphere without Split Edges activated on the Explode modifier.



Uvsphere with Split Edges activated on the Explode modifier.

Unborn

This button controls whether mesh faces will be visible or not before a particle for it has been created/born.

It means that, if you have for example an Explode modifier on a mesh and certain faces on that mesh do not currently have any created/born particles on them, then those faces will not be visible unless the Unborn button is selected, making faces appear to pop into visibility when particles are created/born.

Alive

This button controls whether mesh faces will be visible or not while a particle for it is alive/active.

It means that, if for example you have an Explode modifier on a mesh and certain faces on that mesh have particles that are alive/active on them, then those faces will be visible only if the Alive button is selected.

Dead

This button controls whether mesh faces will be visible or not when the particle associated with the face is dead.

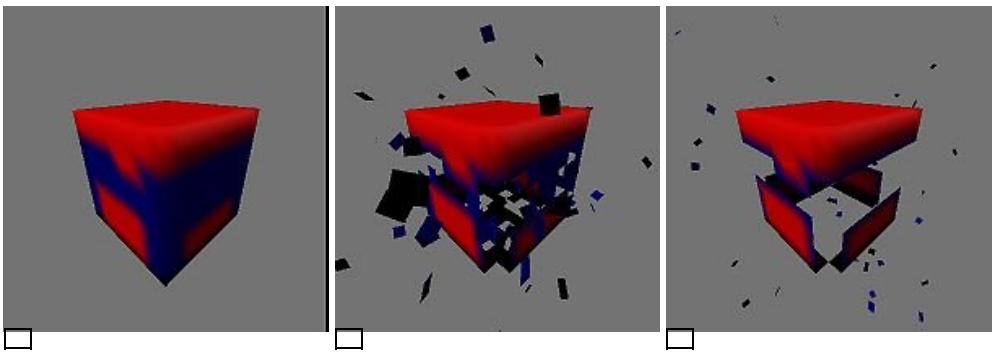
It means that, if for example you have an Explode modifier on a mesh and certain faces on that mesh have particles that are dead on them, then those faces will be visible only if the Dead button is selected.

Examples

Exploding Cube

Weight painted object with *SplitEdge* option not selected

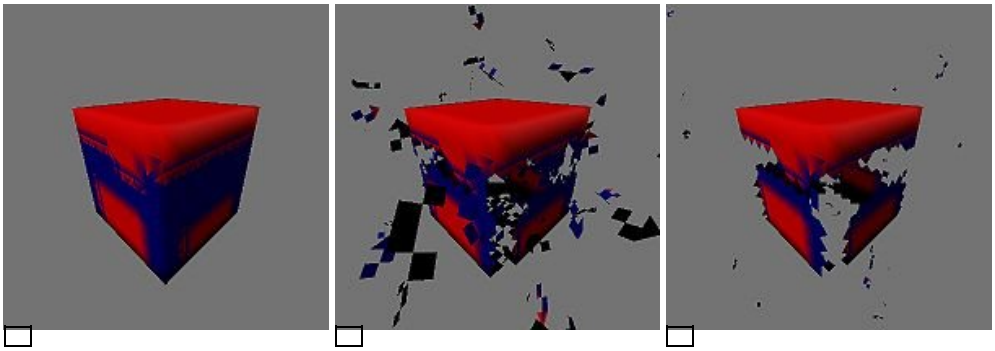
Note that in the image to the right, the faces are not jagged.



Start of simulation. Middle of simulation. End of simulation.

Weight painted object with *SplitEdge* option selected

Note that in the image to the right, the edges are jagged.



Start of simulation. Middle of simulation. End of simulation.

Download .blend file

Media:Manual_-_Explode_Modifier_-_Example_Exploding_Cube.blendâ←

Another explosion

Using the Explode modifier with a particle system to break meshes apart:

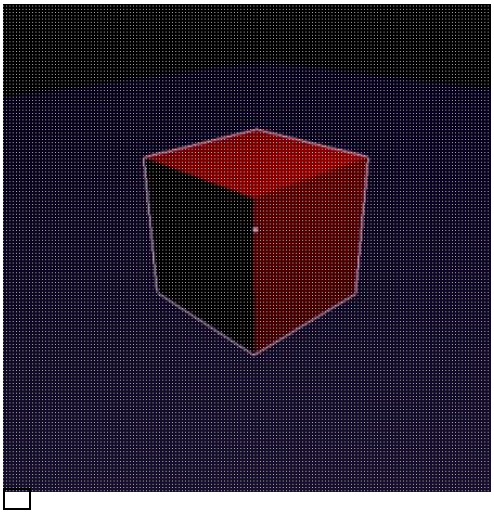
Mask Modifier

Mode: Any mode

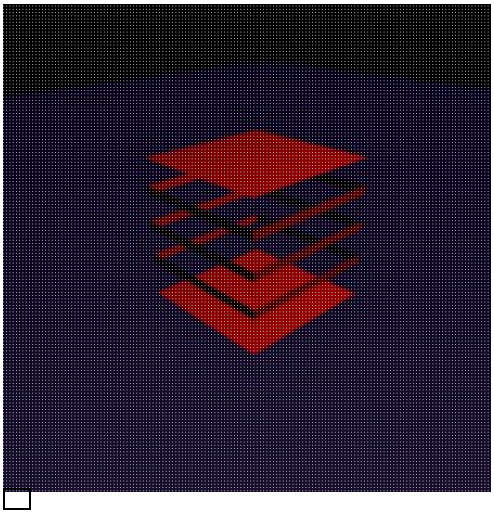
Panel: Modifiers (Editing context, F9)

Description

The Mask modifier allows certain parts of an objects mesh to be hidden from view (masked off), in effect making the parts of the mesh that are masked act as if they were no longer there.



Cube before a Mask modifier is added to it.



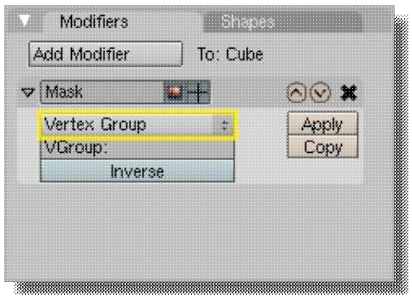
The same cube with the Mask modifier added, with assigned vertex groups.

The screenshots above show a cube before and after a Mask modifier is added to it. In the second screenshot, the cube has had parts of its mesh assigned to a vertex group which the Mask modifier uses to determine which parts will be visible and which parts wonâ← t.

Options

Mask Mode

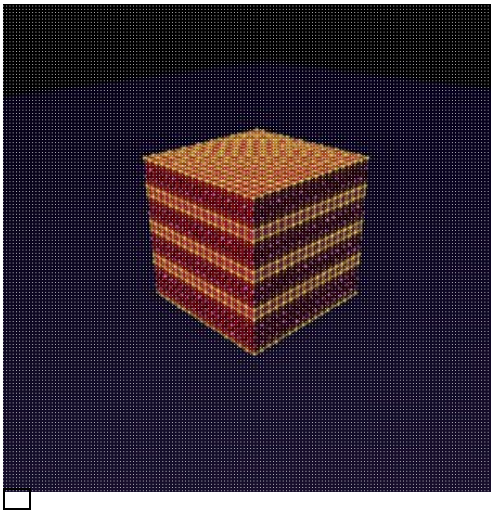
The Mask modifier can hide parts of a mesh based on two different modes, selectable from this drop-down list:



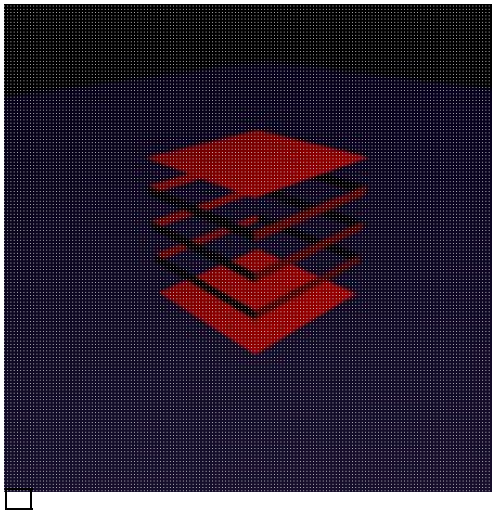
The Mask modifier panel settings with the Vertex Group option highlighted in yellow.

Vertex Group

When the Vertex Group option is selected, the Mask modifier uses a vertex group, specified in the VGroup text field that shows up, to determine which parts of the mesh are masked by the modifier.
Once you have entered the desired name, the Mask modifier will update so that anywhere the vertices of the mesh are part of the named vertex group will be masked (which normally means they will be visibleâ←¦), and anything not part of the named vertex group will be made non-visible.

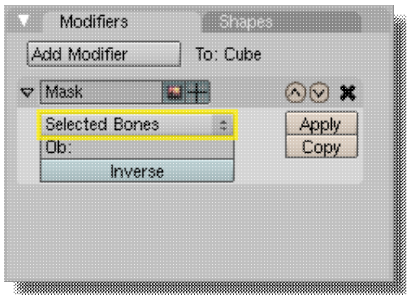


Cube showing the masked parts of the mesh that are part of a selected vertex group (here highlighted in yellow).



The effect of the vertex group on the underlying cube with a Mask modifier active.

The two screenshots above show, first the underlying vertex groups assigned to the cube (highlighted in yellow), and then the effect on the cube when it has a Mask modifier applied to it.
Any of the methods for assigning vertex weights to a mesh work with the Mask modifier, however the actual weight value assigned to a vertex group is completely ignored. The Mask modifier only takes into account whether a set of vertices are part of a group or not, the weight is not taken into account. So having a vertex group weight of say **0.5** will not make a partially masked mesh. Just being part of the vertex group is enough for the Mask modifier, even if the weight is **0.0**.

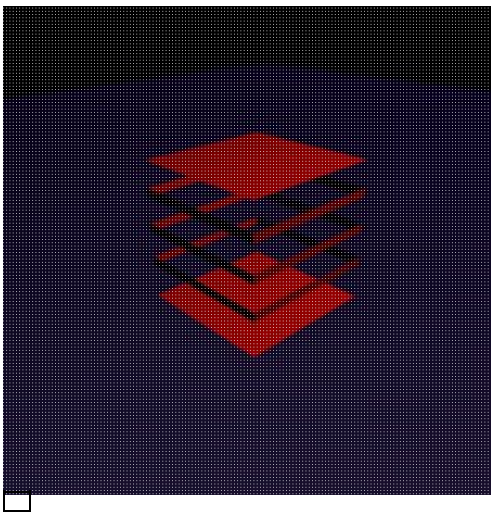


The Mask modifier panel settings with the Selected Bones option highlighted in yellow.

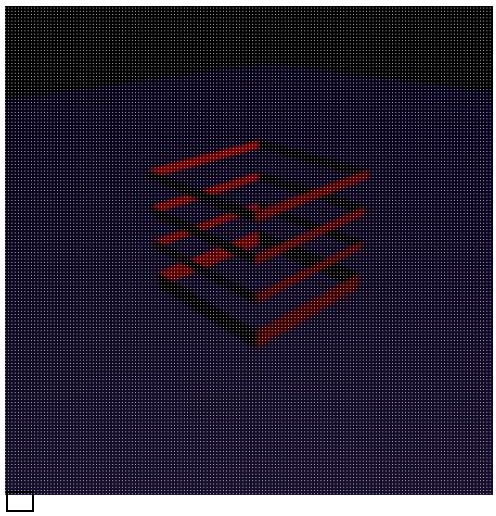
Selected Bones
: Selected Bones is useful in Pose mode or when editing an armature. Enter the name of the armature object in the Ob field that appears. When working with bones in Pose mode, vertex groups not associated with the active bone are masked. The Inverse button can be useful to see how a bone affects the mesh down the chain of bones.

Inverse

Normally, when the Mask modifier is applied to areas of a mesh, the parts that are under the influence of the modifier are left visible while the parts that arenâ ← t are hidden. The Inverse button reverses this behavior, in that now parts of the mesh that were not originally visible become visible, and the parts that were visible become hidden!



A cube with a Mask modifier, without the Inverse option.



A cube with a Mask modifier, with the Inverse option active.

Examples

Example Blend file using Mask modifier

Mirror Modifier

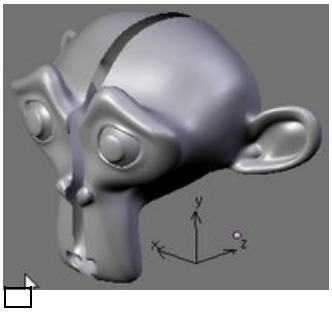
Mode: Any mode

Panel: Modifiers (Editing context, F9)

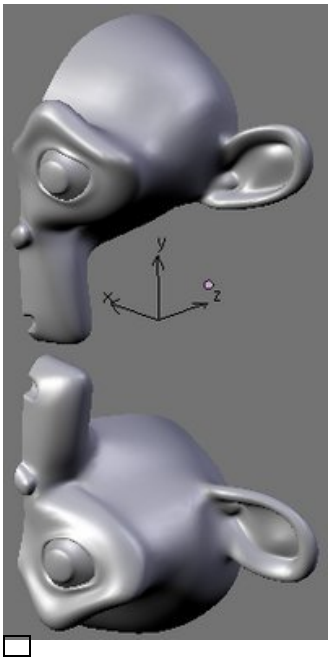
Description

The Mirror modifier automatically mirrors a mesh along its **local** X, Y and/or Z axes, which pass through the objectâ← s center (the mirror plane is then defined by the two other axes). It can also use another object as mirror center, then using this objectâ← s local axes instead of its own. It can weld vertices together on the mirror plane within a specified *tolerance* distance. Vertices from the original object can be prevented from moving across or through the mirror plane. And last but not least, it can also mirror vertex groups and UV coordinatesâ←¦

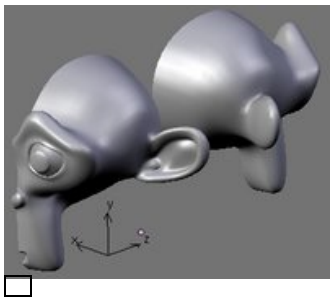
The examples below show the mirror modifier in action in all three axes. The object center, pink dot, has been moved from the â← realâ← center so that the mirror plane is more obvious to spot.



Mirror X axis.

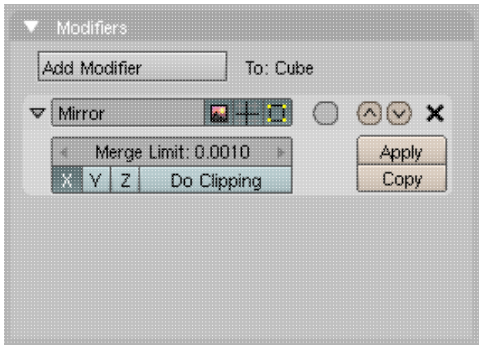


Mirror Y axis.



Mirror Z axis.

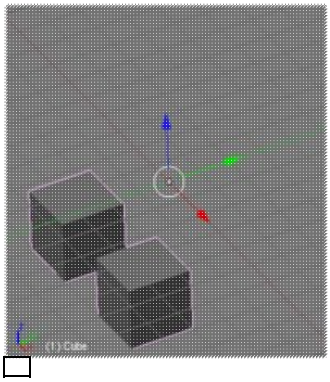
Options

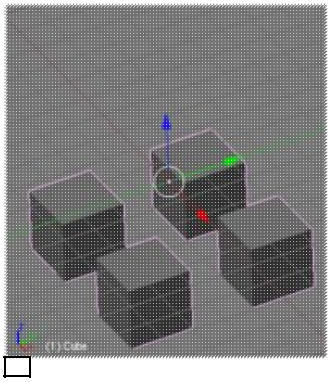


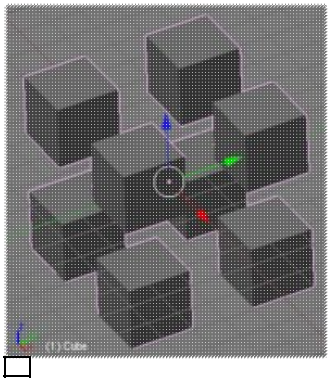
Mirror modifier.

Merge Limit
: The maximal distance between vertices and mirror plane for the welding between original and mirrored vertices to take place. The vertices then will snap together, allowing to link the original mesh to its mirrored copy.

X, Y, Z
: The axis along which to mirror (i.e. the axis perpendicular to the mirror plane of symmetry). To understand how the axis applies to the mirror direction, if you where to mirror on the X axis, the X plus values of the original mesh would become X minus values on the mirrored instance.
You can select more than one of these axes â← Â youâ← ll then get more mirror instances, so that all planes of symmetry selected are â← fully processedâ← (i.e. with two axes, youâ← ll get four objects, and with the three axes selected, height objectsâ←¦).

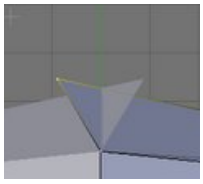

Mirror along X axis.


Mirror along X and Y axes.


Mirror along X, Y and Z axes.

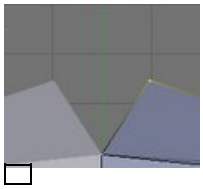
Mirror along several axes. Note that the objectâ← s center has been moved away (the red/green/blue arrows).

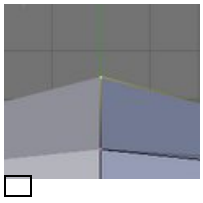
When you use another object (see Ob field, below), it will be its axes that are used as the mirror onesâ←¦



Vertices do not clip together.

Do Clipping
: Prevents vertices from crossing through the mirror plane(s). Note that this is only valid in Edit mode (i.e that when using object transformations, translations, scaling, etc., in Object mode, vertices will happily cross these bordersâ←¦).
While the Do Clipping is un-selected the vertices will not clip together at the mirror plane even if they are within the Merge Limit.





No vertex clipping. Vertices have clipped together.

If Do Clipping is selected but vertices are outside of the Merge Limit the vertices will not merge. As soon as the vertices are within Merge Limit they are clipped together and cannot be moved beyond the mirror plane. If several vertices are selected and are at different distances to the mirror plane, they will one by one be clipped at the mirror plane.
Once you have confirmed clipped vertices with LMB you must, if you want to break the clipping, un-select the Do Clipping to be able to move vertices away from the mirror.

Mirror VGroups

When this button is enabled, the Mirror modifier will try to mirror existing vertex groups. A very nice feature, but that has quite specific prerequisites.

- ◊ First, the vertex groups you want to mirror must be named following the usual left/right pattern (i.e. suffixed by something like â← .Râ← , â← .rightâ← , â← .Lâ← , etcâ←¦).
- ◊ Next, you must have the â← mirroredâ← groups already existing (i.e. same names suffixed by the â← other sideâ←) *and completely empty* (no vertex assigned to it), else it wonâ← t workâ←¦

Usually, the mirrored copies of the vertices of a group remain in this group. Once this option is activated, all groups following the rules described above will only be valid on the original object â← Â the mirrored copy will put these same vertices into the â← mirrorâ← group. Very handy with armatures, for exampleÂ : you just model half of your object, carefully rig it with half of your armature, and just let the Mirror modifier build the other halfâ←¦ Just be sure to put your Armature modifier(s) after the Mirror one!
A final word about multi-axes mirrorÂ : in these cases, the â← directâ← , â← first levelâ← copies get the mirrored groups, the copies of copies (â← second levelâ←) get the original groups, etc.

Mirror U and Mirror V

These two options allows you to mirror, respectively, the U and V texture coordinates. The values are â← mirroredâ← around the **0.5** value, i.e. if you have a vertex with UV coordinates of (**0.3**, **0.85**), its mirror copy will have UV coordinates of (**0.7**, **0.15**) with both buttons enabledâ←¦

Ob

The name of another object (usually an empty), to be used as reference for the mirror processÂ : its center and axes will drive the plane(s) of symmetry. You can of course animate its position/rotation (Ipo curves or others), to animate the mirror effect!

Hints

Many modeling tasks involve creating objects that are symmetrical. However there used to be no quick way to model both halves of an object without using one of the workarounds that have been discovered by clever Blender artists over the years.

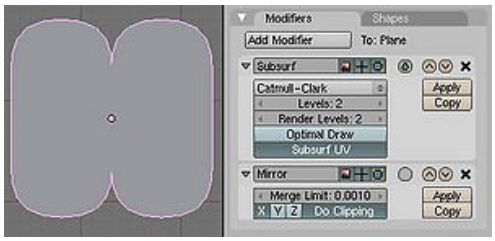
A common technique is to model one half of an object and use AltD to create a linked duplicate which can then be mirrored on one axis to produce a perfect mirror-image copy, which updates in realtime as you edit.

The Mirror modifier offers another, simpler way to do this. Once your modeling is completed you can either click Apply to make a real version of your mesh or leave it as is for future editing.

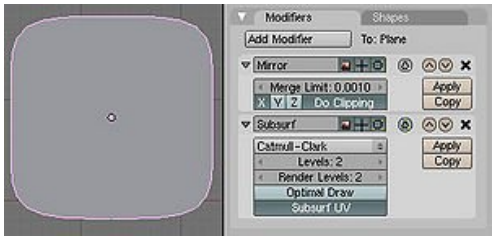
Mirror modeling is actually done in Edit mode.

Using Mirror modifier with Subsurf modifier

When using the Mirror modifier along with the Subsurf modifier, the order in which the modifiers are placed is important.



Subsurf modifier before Mirror modifier.
This shows the Subsurf modifier placed before the Mirror one, as you can see the effect of this is that the mesh splits down the center line of the mirror effect.



Mirror modifier before Subsurf modifier.
This shows the Mirror modifier placed before the Subsurf modifier. In this order you will get the the center line of the mesh snapped to the center line, which in most cases would be the desired effect.

Aligning for Mirror

To apply a Mirror modifier, it is common to have to move the objectâ← s center onto the edge or face that is to be the axis for mirroring. This can be tricky when attempted visually. A good technique to achieve an exact position is to determine the edge against which you wish to mirror. Select two vertices on that edge. Then use â § ShiftS followed by CursorÂ ->Â Selection. This will center the 3D cursor exactly on the edge midway between the two vertices. Finally, in the Editing context (F9), select Center Cursor from the Mesh panel, which will move the objectâ← s center to where the 3D cursor is located. Now the mirroring will be exact.

Or just use another empty as mirror center (Ob field)Â !

Smooth Modifier

Mode: Any mode

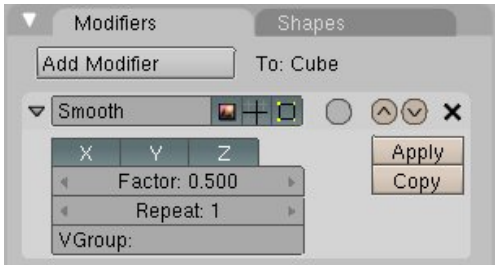
Panel: Modifiers (Editing context, F9)

Description

This modifier smooths a mesh by flattening the angles between adjacent faces in it, just like the Smooth button in the Editing context, Mesh Tools panel. So it smooths without subdividing the mesh â← Â the number of vertices remains the same.

This modifier is not limited to smoothing, though. Its control factor can be configured outside the [**0.0**, **1.0**] range (including negative values), which can result in interesting deformations, depending on the affected mesh.

Options



Smooth modifier.

X, Y, Z
: Toggle buttons to enable/disable the modifier in the X, Y and/or Z axes directions.

Factor
: The factor to control smoothing amount. The smoothing ranges from **0.0** to **1.0** (**0.0**: disabled, **0.5**: same as the Smooth button, **1.0**: maximum). Alternatively, values outside this range (above **1.0** or below **0.0**) distort the mesh.

Repeat
: The number of smoothing iterations, equivalent to pressing the Smooth button multiple times.

VGroup
: A vertex group name, to restrict the effect to the vertices in it only. This allows for selective, realtime smoothing, by painting vertex weights.

Examples



Smooth modifier examples.
Repeat option is set to **5** for all models, but each one has a different Factor value. From left to right:
Top (smoothing): 1) normal Suzanne (same as Factor: **0.0**); 2) Factor: **0.5**; 3) **1.0**; 4) **2.0**.
Bottom (distortion): 1) Factor: **-0.3**; 2) **-0.6**; 3) **-1.0**.
The three models at the bottom row were subdivided twice to look nicer.

Subdivision Surfaces

Mode: Any mode

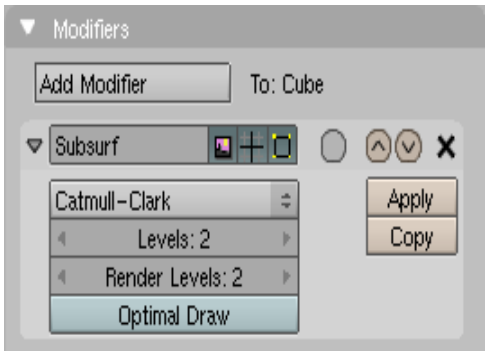
Panel: Modifiers (Editing context, F9)

Hotkey: â § ShiftO (toggle subsurf in Object mode)

Description

Subdivision Surface is a method of subdividing the faces of a mesh to give a smooth appearance, to enable modeling of complex smooth surfaces with simple, low-vertex meshes. This allows high resolution mesh modeling without the need to save and maintain huge amounts of data and gives a smooth *organic* look to the object. With any regular mesh as a starting point, Blender can calculate a smooth subdivision on the fly, while modeling or while rendering, using a simple dummy subdivision surface (Subsurf in short), or the smarter Catmull-Clark one.

Options



Modifiers panel.
Subsurf is a modifier. To add it to a mesh, press Add Modifier and select Subsurf from the list.

Subsurf Type
This drop-down list allows you to choose the subdivision algorithm:
- ◊ Simple Subdiv. â← Just subdivides the surfaces, without any smoothing (similar to Levels time W â Subdivide, in Edit mode). Rarely useful!
- ◊ Catmull-Clark â← Default option, subdivides and smooths the surfaces.

Levels
Defines the display resolution, or level of subdivision *for Blender 3D views*.

Render Levels
This is the subdivision level *used during rendering*.

Note
These two settings allow you to keep a fast and lightweight approximation of your model when interacting with it in 3D, but use a higher quality version when rendering.

Optimal Draw

Restricts the wireframe display to only show the original mesh cage edges, rather than the subdivided result, to help visualization.

Subsurf UV
: When enabled, the UV maps will also be subsurfed (i.e. Blender will add â← virtualâ← coordinates for all sub-faces created by this modifier).

To view and edit the results of the subdivision (â← isolinesâ←) while youâ← re editing the mesh, you must enable the Editing Cage mode by clicking in the gray circle in the modifier panel header (next to the arrows for moving the modifier up and down the stack). This lets you grab the points as they lie in their new subdivided locations, rather than on the original mesh.

Hints

You can use â § ShiftO if you are in Object mode to switch Subsurf On or Off in *both* the Object mode *and* at rendering. To only turn the subsurf view on/off in Object mode (to reduce lag), press â § ShiftAltO. The Subsurf level can also be controlled via Ctrl1 to Ctrl4, but this only affects the visualization sub-division level.

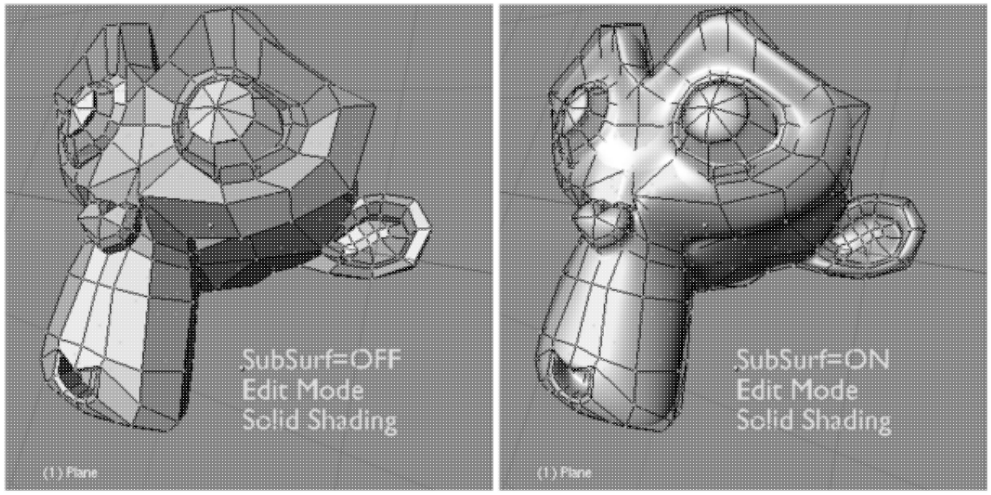
A subsurfed mesh and a NURBS surface have many points in common such as they both rely on a â← coarseâ← low-poly â← meshâ← to define a smooth â← high definitionâ← surface, however, there are notable differences:

- NURBS allow for finer control on the surface, since you can set â← weightsâ← independently on each control point of the control mesh. On a subsurfed mesh you cannot act on weights (except for edges crease, see bellow).
- Subsurfs have a more flexible modeling approach. Since a subsurf is a mathematical operation occurring on a mesh, you can use all the modeling tools and techniques described in the chapter on mesh modeling. There are more techniques, which are far more flexible, than those available for NURBS control polygons.

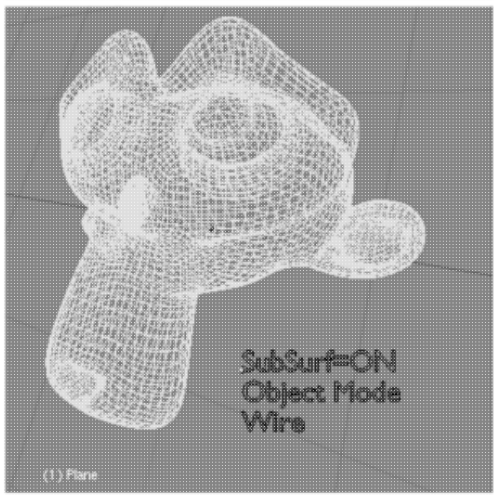
Since Subsurf computations are performed both real-time, while you model, and at render time, and they are CPU intensive, it is usually good practice to keep the subsurf Levels low (but non-zero) while modeling, and higher while rendering (Render Levels).

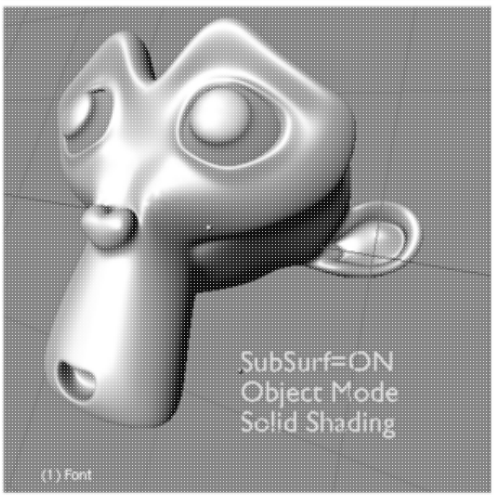
Examples

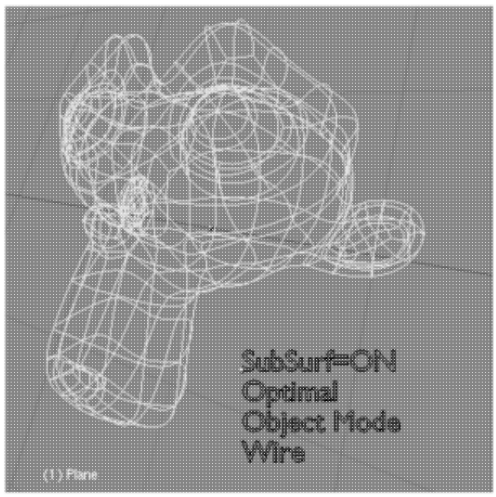
(*Subsurfed Suzanne*) is a series of pictures showing various different combinations of Subsurf options on a Suzanne mesh.

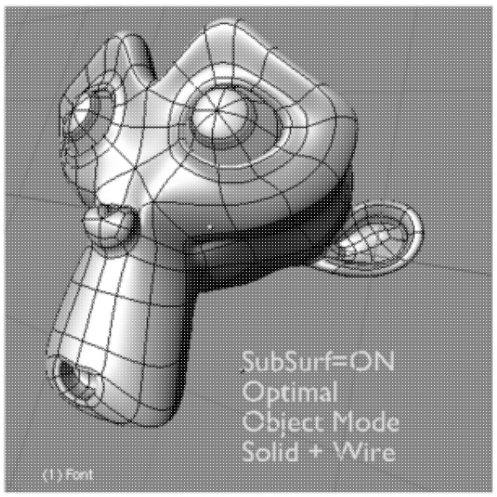

SubSurf=OFF
Edit Mode
Solid Shading
(1) Plane

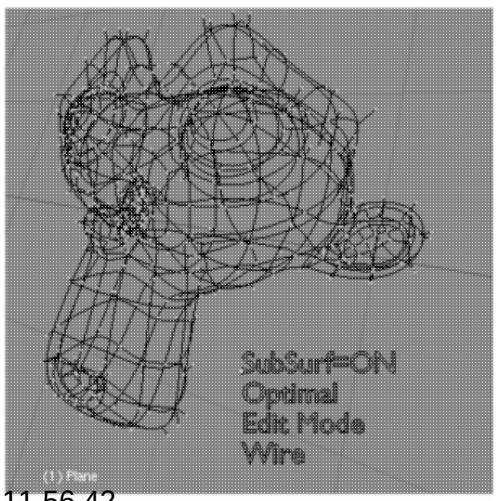
SubSurf=ON
Edit Mode
Solid Shading
(1) Plane

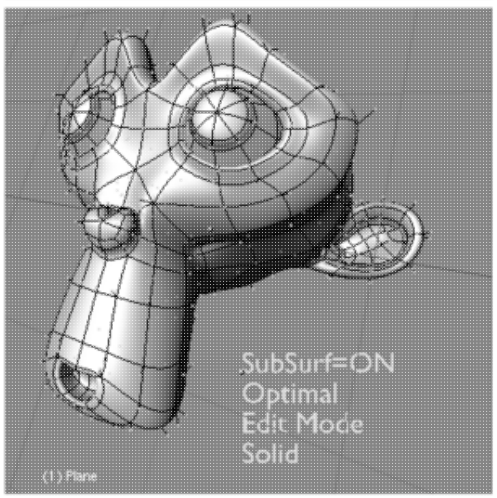

SubSurf=ON
Object Mode
Wire
(1) Plane


SubSurf=ON
Object Mode
Solid Shading
(1) Font


SubSurf=ON
Optimal
Object Mode
Wire
(1) Plane


SubSurf=ON
Optimal
Object Mode
Solid + Wire
(1) Font


SubSurf=ON
Optimal
Edit Mode
Wire
(1) Plane


SubSurf=ON
Optimal
Edit Mode
Solid
(1) Plane





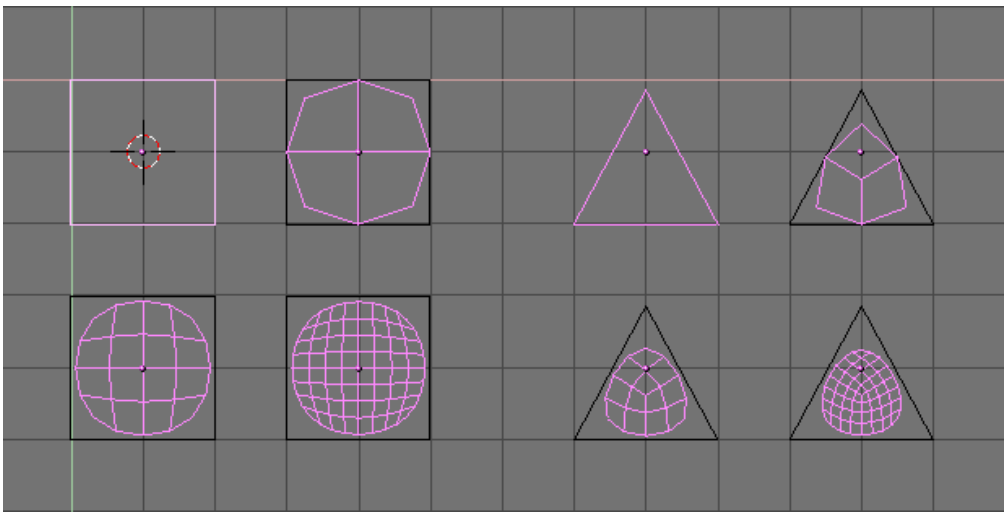
Subsurfed Suzanne.

(*SubSurf of simple square and triangular faces*) shows a **0**, **1**, **2**, **3** level of subsurf on a single square face or on a single triangular face. Such a subdivision is performed, on a generic mesh, for *each* square or triangular face.

It is obvious that the number of faces produced by the n level of subdivision is:

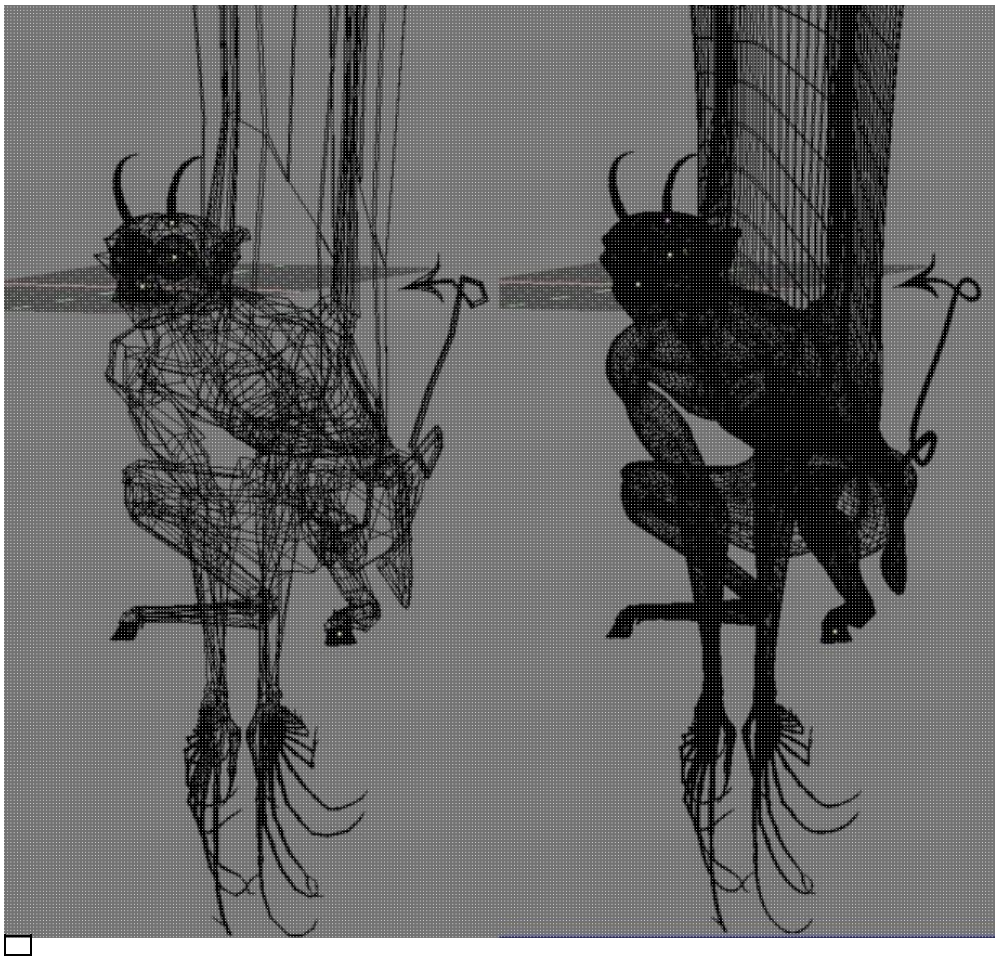
- For a quadrangle : 4^n.
- For a triangle : 3 Ã 4^{n-1}.

This dramatic increase of face (and vertex) number results in a slow-down of all editing, and rendering, actions and calls for lower Subsurf level in the editing process than in the rendering one.

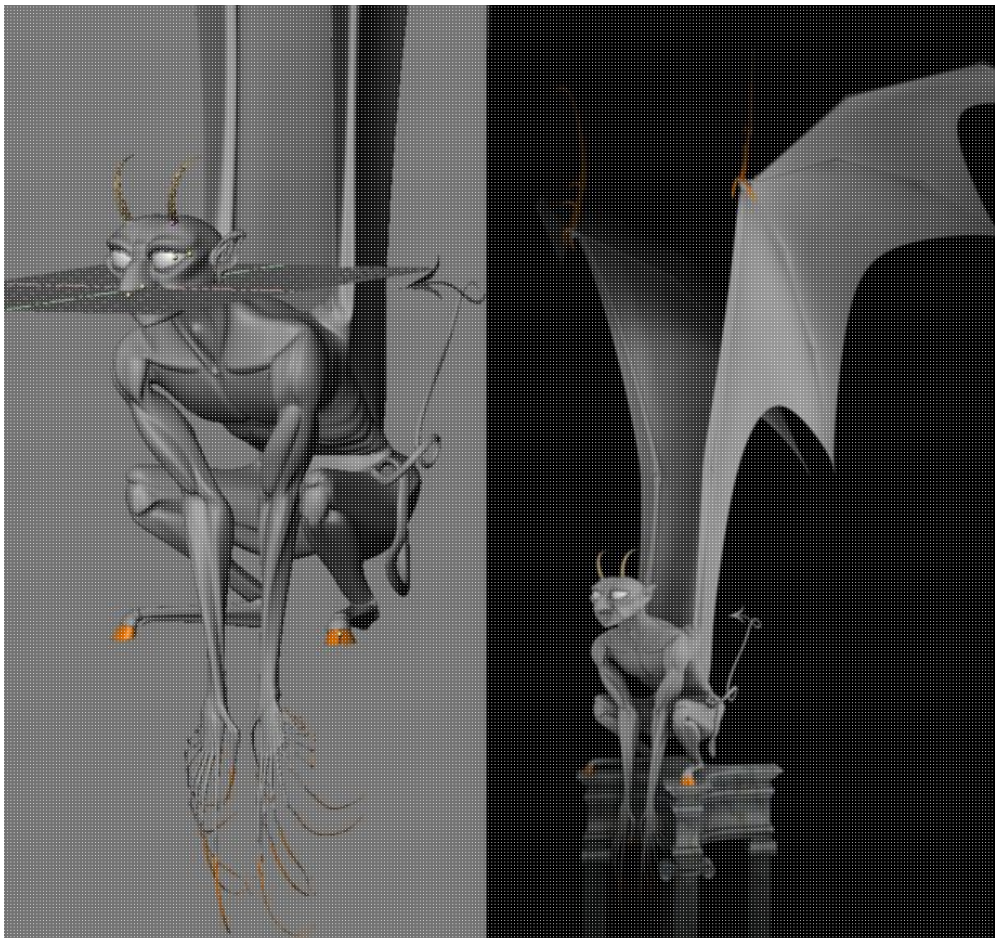


SubSurf of simple square and triangular faces.

The Subsurf tool allows you to create very good â← organicâ← models, but remember that a regular mesh with square faces, rather than triangular ones, gives the best results. (*A Gargoyle base mesh (left) and pertinent level **2** subsurfed mesh (right)*) and (*Solid view (left) and final rendering (right) of the Gargoyle*) show an example of what can be done with Blender subsurfs.



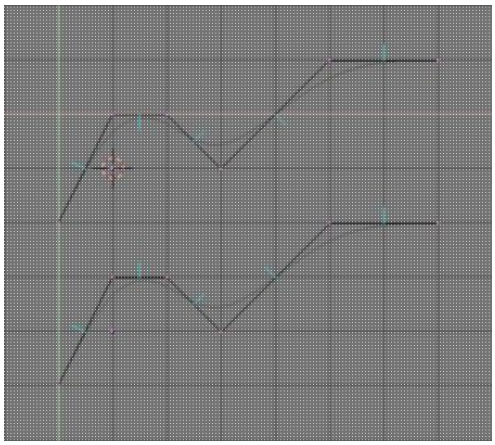
A Gargoyle base mesh (left) and pertinent level **2** subsurfed mesh (right).



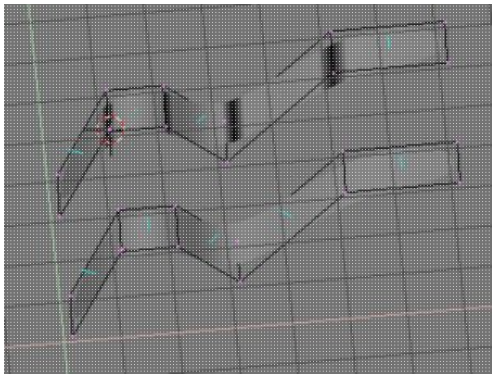
Solid view (left) and final rendering (right) of the Gargoyle.

Limitations & Workarounds

Blenderâ← s subdivision system is based on the Catmull-Clark algorithm. This produces nice smooth subsurfed meshes but any subsurfed face, that is, any small face created by the algorithm from a single face of the original mesh, shares the normal orientation of that original face.



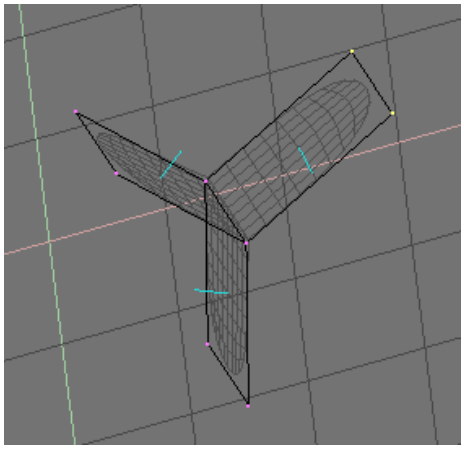
Side view of subsurfed meshes. With random normals (top) and with coherent normals (bottom).
This is not an issue for the shape itself, as (*Side view of subsurfed meshes. With random normals (top) and with coherent normals (bottom)*) shows, but it is an issue in the rendering phase and in solid mode, where abrupt normal changes can produce ugly black lines (*Solid view of subsurfed meshes with inconsistent normals (top) and consistent normals (bottom)*).



Solid view of subsurfed meshes with inconsistent normals (top) and consistent normals (bottom).
Use the CtrlN command in Edit mode, with all vertices selected, to recalculate the normals to point outside.

In these images the face normals are drawn cyan. You can enable drawing normals in the Editing context (F9), Mesh Tools More panel.

Note that Blender cannot recalculate normals correctly if the mesh is not â← Manifoldâ← . A â← Non-Manifoldâ← mesh is a mesh for which an â← outâ← cannot unequivocally be computed. From the Blender point of view, it is a mesh where there are edges belonging to *more* than two faces.



A â← Non-Manifoldâ← mesh.
(*A â← Non-Manifoldâ← mesh*) shows a very simple example of a â← Non-Manifoldâ← mesh. In general a non-manifold mesh occurs when you have internal faces and the like.

A non-manifold mesh is not a problem for conventional meshes, but can give rise to ugly artifacts when subsurfed. Also, it does not allow decimation, so it is better to avoid them as much as possible.

Use these two hints to tell whether a mesh is â← Non Manifoldâ← :

- The recalculation of normals leaves black lines somewhere.
- The Decimator tool in the Mesh panel refuses to work stating that the mesh is â← Non-Manifoldâ← .

Weighted creases for subdivision surfaces

Mode: Edit mode (mesh)

Panel: 3D View â Transform Properties (N)

Hotkey: â § ShiftE

Menu: Mesh Â» Edges Â» Crease Subsurf

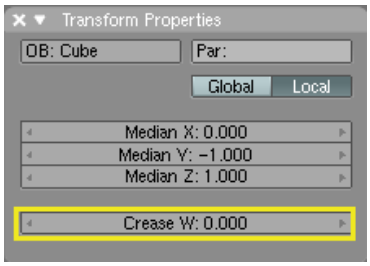
Description

Weighted edge creases for subdivision surfaces allow you to change the way Subsurf subdivides the geometry to give the edges a smooth or sharp appearance.

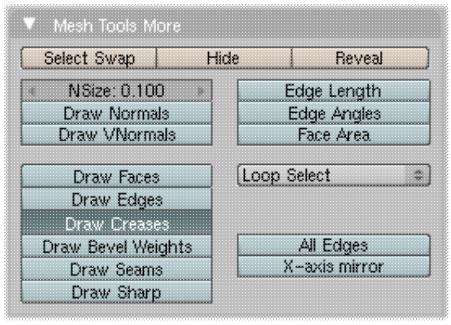
Options

The crease weight of selected edges can be changed interactively by using â § ShiftE and moving the mouse towards or away from the selection. Moving the mouse away from the edge increases the weight. You can also use Transform Properties (N) and enter the value directly. A higher value makes the edge â← strongerâ← and more resistant to subsurf. Another way to remember it is that the weight refers to the edgeâ← s sharpness. Edges with a higher weight will be deformed less by subsurf. Recall that the subsurfed shape is a product of all intersecting edges, so to make the edges of an area sharper, you have to increase the weight of all the surrounding edges.

You can enable an indication of your edge sharpness by enabling Draw Creases. See (*Mesh Tools More panel*).



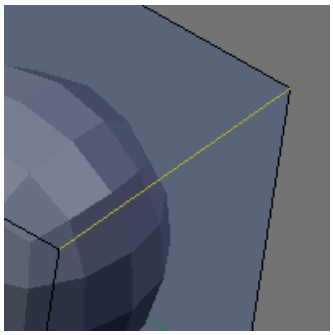
Transform Properties panel.



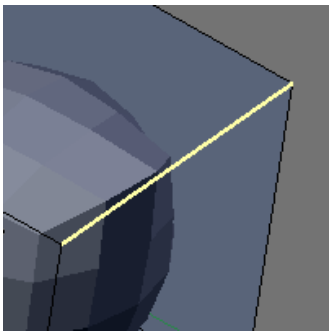
Mesh Tools More panel.

Examples

The sharpness value on the edge is indicated as a variation of the thickness on the edge. If the edge has a sharpness value of **1.0**, it will be thicker than if it had a sharpness value of **0.0**.



Crease: **0.0**.



Crease: **1.0**.

UVProject Modifier

Mode: Any mode

Panel: Modifiers (Editing context, F9)

Description

Ok, so have you ever seen someone walk in front of the slide projector, and the slide is projected on them as they walk in front of it? Well, this modifier does that to the object, using an image (still, movie clip, etc) that you specify.

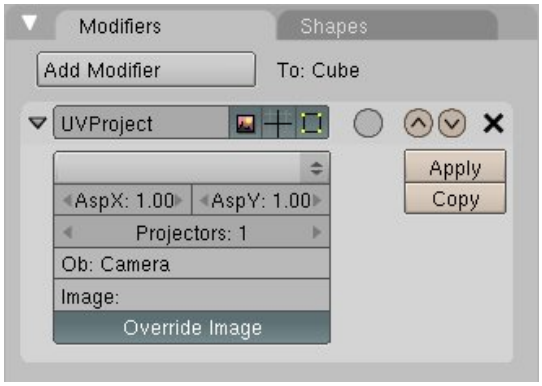
The UVProject modifier projects an UV-mapped image onto the base mesh as if from a slide projector. It does this by altering the base meshâ ← s UV coordinates as though they were projected onto the object by one or more projector objects. The projection used is either an orthogonal or perspective projection along the projector objectâ ← s negative Z axis.

The effect can be limited to just those faces associated with a given image, or all faces can be affected. The modifier can also override the image assigned to each face with a given image.

For the results to be visible in the final render, the base mesh must have:

- A material that uses the affected UV coordinate layer in some way (e.g. by having TexFace enabled).
- The named UV texture active.

Options



UVProject modifier.

UV Layer
: Selects the UV coordinate layer to affect. If no UV layers are present, the modifier will be disabled.

AspX, AspY
: Sets the aspect ratio of the image being projected.

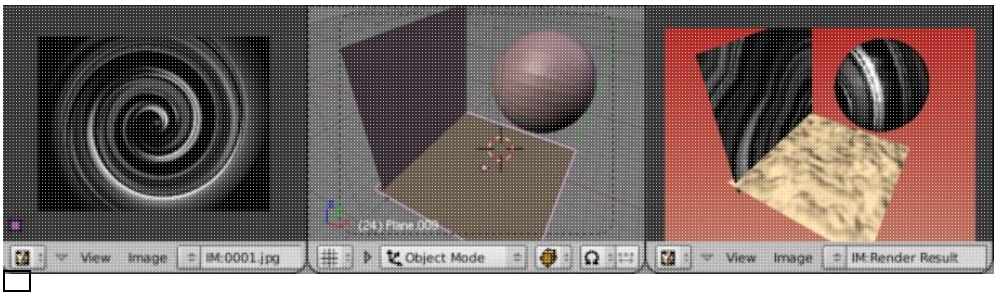
Projectors
: The number of projector objects to use (from **1** to **10**). An Ob field will be added to the panel for each projector.

Ob
: A projector object to use. There will be from one to ten of these fields, depending on the value of Projectors. For each face, the projector whose projection axis is most perpendicular to that face will be used.
 - ◊ If a projector object is a camera, the projection used will depend on the camera type: orthographic projection for orthographic cameras and perspective projection for perspective cameras.
 - ◊ If a projector object is not a camera, the projection will always be orthographic.

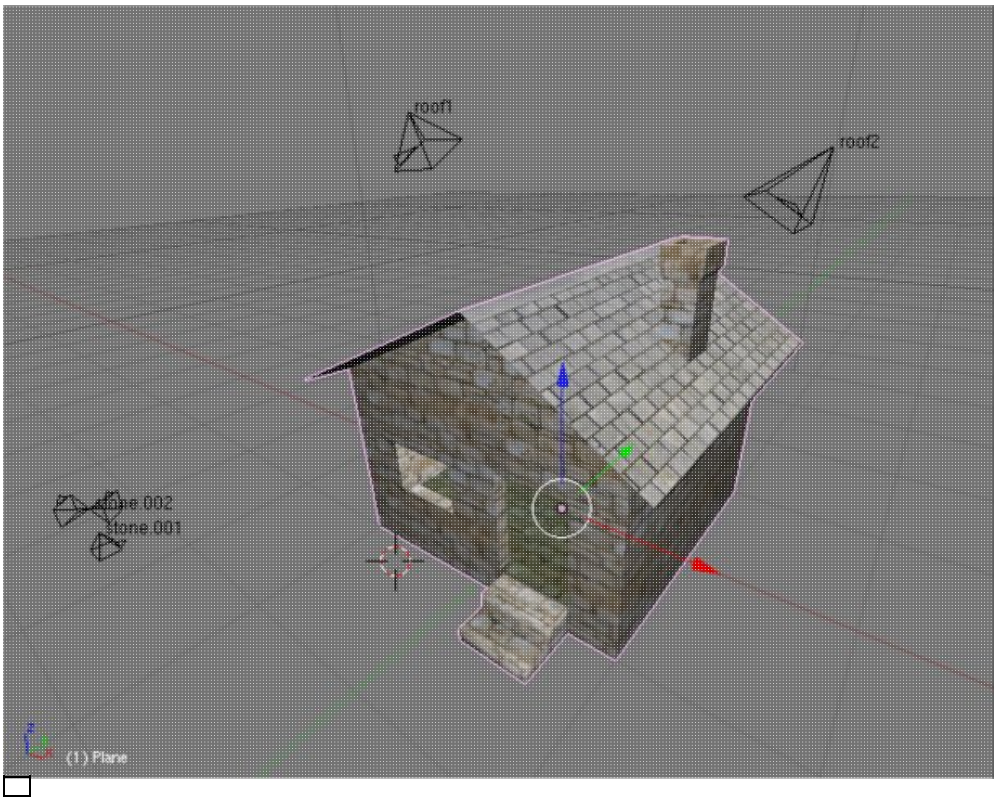
Image
: The image to use.
 - ◊ If Image is empty, **all faces** are affected.
 - ◊ If Image is not empty, **only faces with that image assigned** are affected (unless Override Image is on, see below).

Override Image
: This button controls whether to override the current face image with the contents of Image. If Override Image is on, **all faces** will be affected by the modifier, and will have the contents of Image assigned to them (if Image is blank, faces will have no image assigned).

Examples



A swirl image is loaded in the UV/Image Editor and thus known to Blender. A simple scene, shown in the middle, has a wall and a sphere with the modifier set to project the swirl from the 35mm (perspective) camera onto those two objects. The two objects have the same material as the floor, but their colors are superseded by the modifier. As the sphere moves (not shown), the part of the image that it takes moves correspondingly. Kinda spooky, if you ask me.



A house textured with the UVProject modifier. The cameras scattered around the scene are the projection objects. UVProject greatly simplifies this kind of texturing.
Sample blend file

Introduction

Lighting is a very important topic in rendering, standing equal to modeling, materials and textures. The most accurately modeled and textured scene will yield poor results without a proper lighting scheme, while a simple model can become very realistic if skillfully lit. Lighting, sadly, is often overlooked by the inexperienced artist who commonly believes that, since real world scenes are often lit by a single light (a lamp, the sun, etc.) a single light will also do in computer graphics. This is false because in the real world even if a single light source is present, the light shed by such a source bounces off objects and is re-irradiated all over the scene, making shadows soft and shadowed regions not pitch black, but partially lit.

Viewing Restrictions

The colour of an object and the lighting of your scene is affected by:

- Your ability to see different colours (partial colour blindness is common).
- The medium in which you are viewing the image (e.g. an LCD panel versus printed glossy paper).
- The quality of the image (e.g. a JPEG at **0.4** compression versus **1.0**).
- The environment in which you are viewing the image (e.g. a CRT monitor with glare versus in a dark room, or in a sunshiny blue room).
- Your brainâ ← s perception of the colour and intensity relative to those objects around it and the world background colour.

So, the exact same image viewed by Person A on monitor B in room C may look very different to Person D viewing a printout E of the image while on the subway F.

Global Influences

In Blender, the things under your control that affect lighting are:

- The colour of the world ambient light.
- The use of Ambient Occlusion as a way to cast that ambient light onto the object.
- The degree to which the ambient light colours the material of the object.
- The use of Radiosity, where the colour of one object radiates onto another.
- The render engine used (Blender Internal versus Yafray).
- The lamps in your scene.

The physics of light bouncing around in the real-world is simulated by Ambient Occlusion (a world setting), buffer shadows (which approximate shadows being cast by objects), ray tracing (which traces the path of photons from a light source). Also, within Blender you can use the Radiosity engine. Ray tracing, ambient occlusion, and radiosity are compute-intensive processes. Blender can perform much faster rendering with its internal scan line renderer, which is a very good scan line renderer indeed. This kind of rendering engine is much faster since it does not try to simulate the real behaviour of light, assuming many simplifying hypotheses.

Lighting Settings

Only after the above global influences do you get into adding on light from lamps in your scene. The main things under your control are the:

- Type of light used (Sun, Spot, Lamp, Hemi, etc).
- Colour of the light.
- Position of the light and its direction.
- Settings for each of those lights, including energy and falloff.

Then you are back to how that materialâ← s shader reacts to the light.

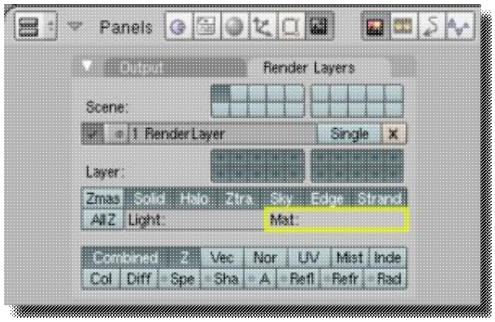
This chapter attempts to address the above, including how lights can work together in rigs to light your scene. In this chapter we will analyze the different type of lights in Blender and their behaviour, we will analyze their strong and weak points. We also describe many lighting rigs, including the ever-popular three point light method.

Lighting in the Workflow

In this User Manual we have placed Lighting before Materials; you should set up your lighting before assigning materials to your meshes. Since the material shaders react to light, without proper lighting, the material shaders will not look right, and you will end up fighting the shader, when it is really the bad lighting that is causing you grief. All of the example images in this section do not use any material setting at all on the ball, cube or background.

Over-riding Materials to reset lighting

If you have started down the road of assigning materials, and are just now fiddling with the lighting, we suggest that you create a default, generic grey material; no VCol, no TexFace, no Shadeless, just plain old middle grey with an RGB of (**0.8**,Â **0.8**,Â **0.8**). If you click the auto-namer button, it should fill in â← `Grey`â← .



Mat field highlighted in yellow.

Next go to Scene context (F10) of the Buttons window, and find the Render sub-context. In the Render Layers panel, enter â← `Grey`â← into the Mat field. If the name sticks, you know you entered it correctly. This will override any materials you may have set, and render everything with this flat boring colour. Using this material, you can now go about adjusting the lighting. Just empty this field to get back to your original materials.

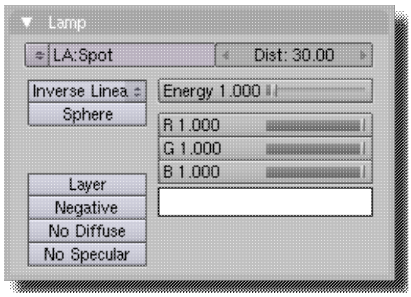
Lights

As previously said, there are multiple types of lighting in Blender, like radiosity or ambient light. However, one of the more used are â← lightsâ← , or â← lampsâ← . In this section, we will talk of general info and settings for these lights (you will find more lamp-specific details in the Lamps part):

- Light Properties â← settings common to all lamps.
- Light Attenuation.
- Textures â← how to apply texture(s) to lamps.
- What Light Affects.
- Lights In Other Contexts â← lamp-related setting in other contexts.

Lights Common Options

Here are some options available for all types of lamps.



Lamp panel.

Dist
: The Dist field indicates the number of Blender Units (BU) at which the intensity of the current light source will be half of its Intensity. Objects less than the number of BU away from the lamp will get more light, while objects further away will receive less light. Certain settings and lamp falloff types affect how the Dist field is interpreted, meaning that it will not always react the same, see the page about light falloff.

Energy
: The Intensity of the light sources illumination (from **0.0** to **10.0**).

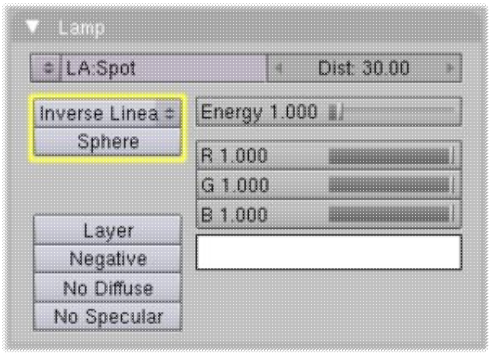
Color
: The color of the light sources illumination.

Light Attenuation

Mode: All modes

Panel: Shading context â Lamp sub-context (F5)

Description



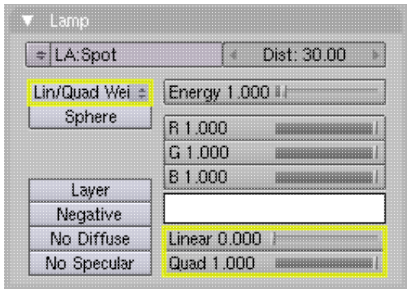
Lamp panel with lamp falloff main controls highlighted in yellow.
The light falloff possibilities has been improved in Blender 2.46, at least for the Lamp and Spot lamps.

There are now two main controls:

- The lamp Falloff types drop-down list.
- The Sphere button

Falloff types

Lin/Quad Weighted



Lamp panel with Lin/Quad Weighted Falloff type selected, and the Quad and Linear sliders also highlighted in yellow.
When this setting is chosen, two sliders are shown, Linear and Quad (previously Quad1 and Quad2), which control respectively the â← linearnessâ← and â← quadraticnessâ← of the falloff curve. Lamps in old files with the Quad button on will be initialized to this setting.

This lamp falloff type is in effect allowing the mixing of the two light attenuation profiles (linear and quadratic attenuation types).

See also the Blender 2.46 release notes.

Linear

This slider input field can have a value between **0.0** and **1.0**. A value of **1.0** in the Linear field and **0.0** in the Quad field, in effect means that the light from this source is completely linear. Meaning that by the number of Blender Units distance specified in the Dist field, this light sources intensity will be half the value it was when it reaches the number of Blender Units distance specified in the Dist field.

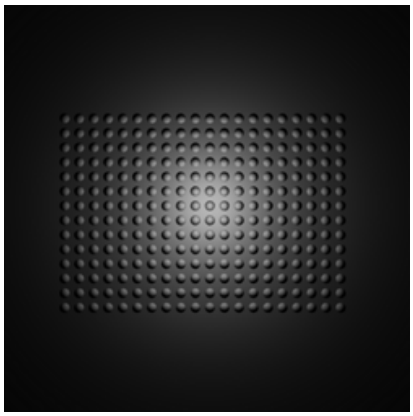
When the Quad slider is set to **0.0**, the formula for working out the attenuation at a particular range for full linear attenuation is:

```
IÂ =Â EÂ Ã Â (DÂ /Â (DÂ +Â LÂ Ã Â r))
```

Where

- `I` is the calculated Intensity of light.
- `E` is the current Energy slider setting.
- `D` is the current setting of the Dist field.
- `L` is the current setting of the Linear slider.
- `r` is the distance from the lamp where the light intensity gets measured.

Quad



Lamp with Lin/Quad Weighted falloff default settings.
This slider input field can have a value between **0.0** and **1.0**. A value of **1.0** in the Quad field and **0.0** in the Linear field means that the light from this source is completely quadratic (quad type).

Quad attenuation type lighting is considered a more accurate representation of how light attenuates (in real world). In fact, fully quadratic attenuation is selected by default on Lin/Quad Weighted lamp fallout (see *Lamp with Lin/Quad Weighted falloff default settings*).

Here again, the light intensity is half when it reaches the Dist distance from the lamp. Comparing the quadratic falloff to the linear falloff, the intensity decays much slowly at distances lower than the set Dist, but it attenuates much quicker after Dist is reached.

When the Linear slider is set to **0.0** the formula for working out the attenuation at a particular range for full quadratic attenuation is:

$$I = E \times D^2 / (D^2 + Q \times r^2)$$

Where

- I is the calculated Intensity of light.
- E is the current Energy slider setting.
- D is the current setting of the Dist field.
- Q is the current setting of the Quad slider.
- r is the distance from the lamp where the light intensity gets measured.

Mixing Linear and Quad

If both the Linear and Quad slider fields have values greater than **0.0**, then the formula used to calculate the light attenuation profile changes to this:

$$I = E \times D / (D + L \times r) \times D^2 / (D^2 + Q \times r^2)$$

Where

- I is the calculated Intensity of light.
- E is the current Energy slider setting.
- D is the current setting of the Dist field.
- L is the current setting of the Linear slider.
- Q is the current setting of the Quad slider.
- r is the distance from the lamp where the light intensity gets measured.

Zeroing both Linear and Quad

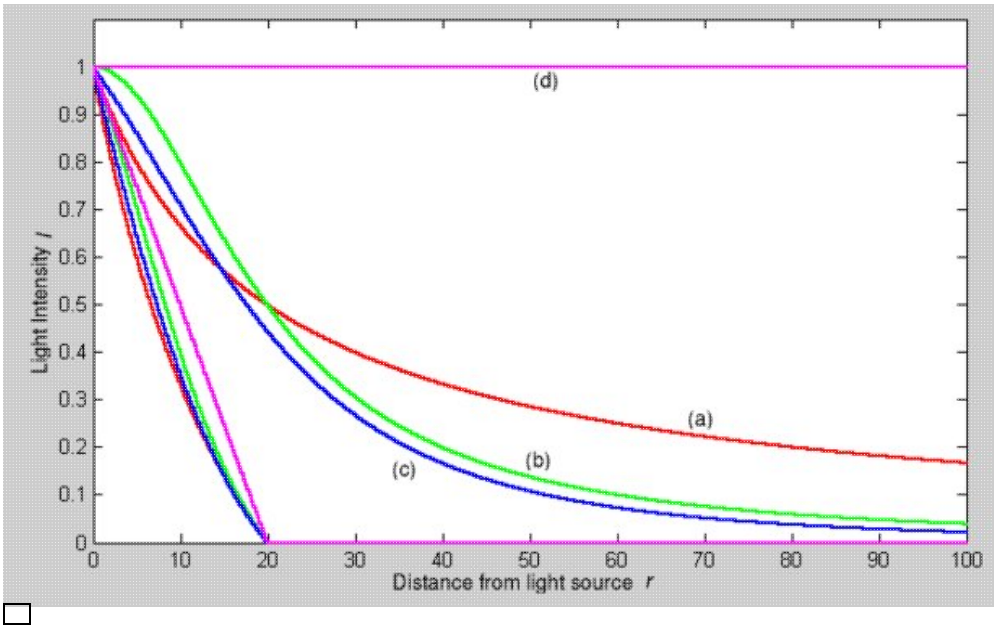
If both the Linear and Quad sliders have **0.0** as their values, the light intensity will not attenuate with distance. This does not mean that the light will not get darker, it will, but only because the energy the light has is spread out over a wider and wider distance. The total amount of energy in the spread out light will remain the same though. Light angle also affects the amount of light you see. It is in fact the behavior of light in the deep space vacuum.

If what you want is a light source that doesn't attenuate and gives the same amount of light intensity to each area it hits, you need a light with properties like the Constant lamp Falloff type.

Also, when the Linear and Quad sliders are both **0.0** values the Dist field ceases to have any influence on the light attenuation, as shown by the equation above.

Graphical Summary

Below is a graph summarizing the lin/quad attenuation type, showing attenuation with or without the Sphere option (described later).



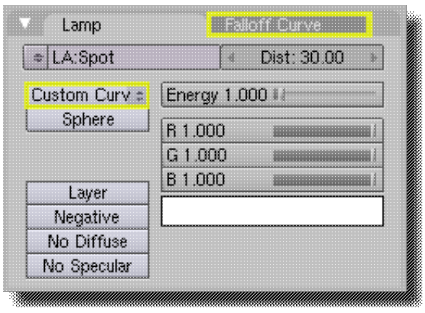
Light Attenuation:
a) Linear (Linear=**1.0**, Quad=**0.0**); b) Quadratic (Linear=**0.0**, Quad=**1.0**);
c) Linear and quadratic (Linear=Quad=**0.5**); d) Null (Linear=Quad=**0.0**).
Also shown in the graph the â← sameâ← curves, in the same colours, but with the Sphere button turned on.

Custom Curve

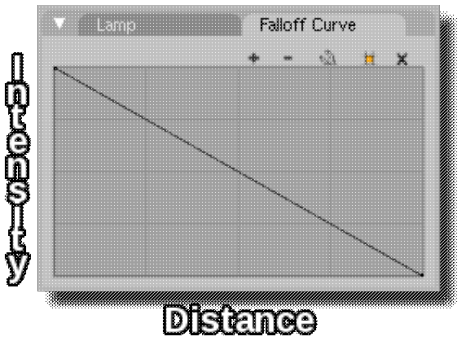
The Custom Curve lamp Falloff type is very flexible.

Most other lamp falloff types work by having their light intensity start at its maximum (when nearest to the light source) and then with some predetermined pattern decrease their light intensity when the distance from the light source gets further away.

When using the Custom Curve Lamp Falloff type, a new panel is created called Falloff Curve. This Falloff Curve profile graph, allows the user to alter how light intense light is at a particular point along a lightâ← s attenuation profile (i.e. at a specific distance from the light source).



Lamp panel with lamp Custom Curve and Falloff Curve tab highlighted.



Falloff Curve panel used to control the amount of attenuation a light has in an arbitrary manner.

The Falloff Curve profile graph has two axes, the â← `Distance`â← axis and the â← `Intensity`â← axis (although these labels were added in the pictures to describe how the Falloff Curve profile graph works, they donâ← t appear in Blender).

Distance axis

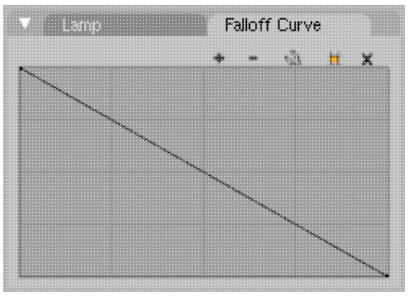
It represents the position at a particular point along a light sources attenuation path. The far left being at the the position of the light source and the far right being the place where the light sources influence would normally be completely attenuated. I say â← normally wouldâ← because the Falloff Curve can be altered to do the exact opposite if required.

Intensity axis

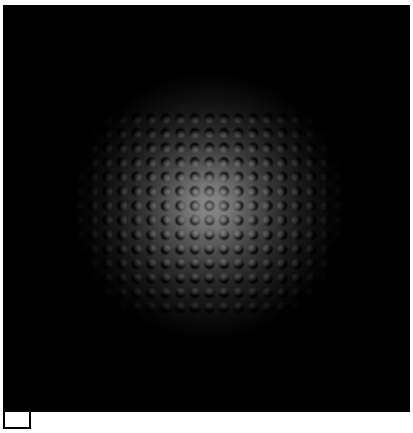
It represents the intensity at a particular point along a light sources attenuation path. Higher intensity is represented by being higher up the intensity axis, while lower intensity light is represented by being lower down on the intensity axis.

Altering the Falloff Curve profile graph is easy. Just LMB click on a part of the graph you want to alter and drag it where you want it to be. If when you click you are over or near one of the tiny black square handles, it will turn white indicating that this is the handle that is now selected, and you will be able to drag it to a new position. If when you click on the graph you are not near a handle, one will be created at the point that you clicked, which you can then drag where you wish.

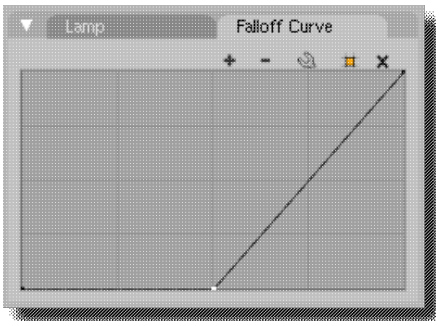
In the example below (the default for the Falloff Curve Profile Graph), the graph shows that the intensity of the light starts off at its maximum (when near the light), and linearly attenuates as it moves to the right (further away from the light source).



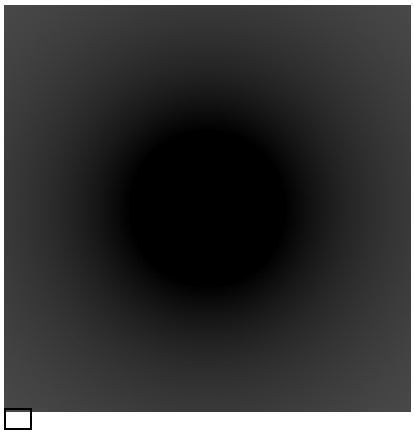
Default Falloff Curve panel graph.


Render showing the Custom Curve lamp falloff type effect with default settings.

If you wanted to have a light attenuation profile that got more intense as it moved away from the light source, you could alter the graph as below:

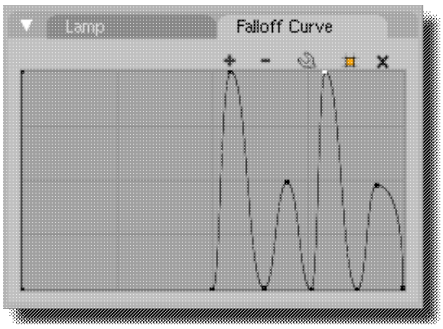


Falloff Curve for reversed attenuation.

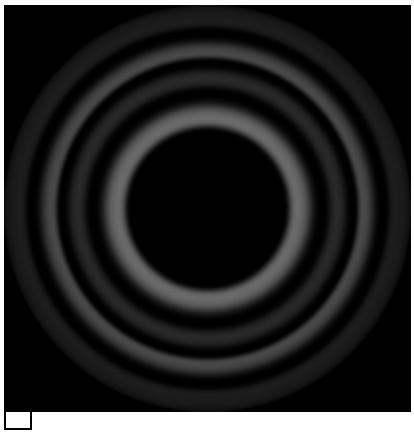

Falloff Curve for reversed attenuation rendered.

You are obviously not just limited to simple changes such as light reversing the attenuation profile, you can have almost any profile you desire.

Here is another example of different Falloff Curve profile graph, along with its resultant render output:

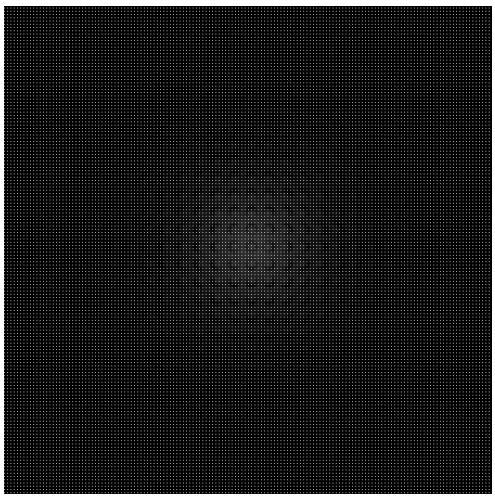


Oscillating attenuation profile.



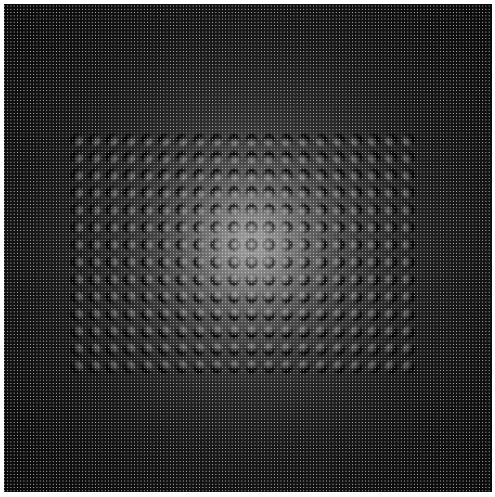
Render showing the effects of a â← waveletâ← profile graph on the light attenuation.

Inverse Square



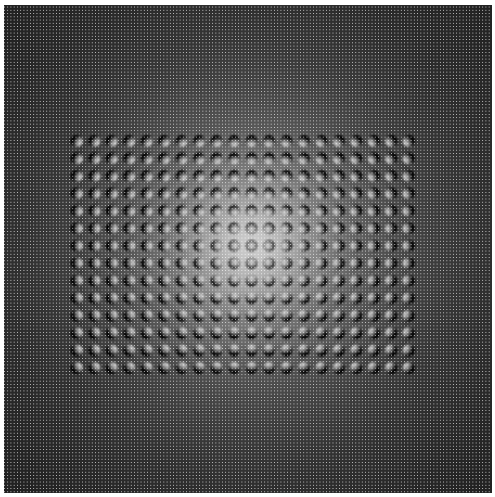
Render showing the Inverse Square lamp falloff type effect with default settings.
This lamp fallout type attenuates its intensity according to inverse square law, scaled by the Dist value. Inverse square is a sharper, realistic decay, useful for lighting such as desk lamps and street lights. This is similar to the old Quad option (and consequently, to the new Lin/Quad Weighted option with Linear to **0.0** and Quad to **1.0**), with slight changes (TODO: Is it true? What changes? -- Mont29 13:42, 20 June 2009 (UTC)).

Inverse Linear



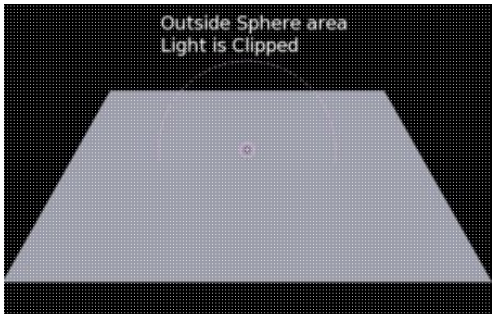
Render showing the Inverse Linear lamp falloff type effect with default settings.
This lamp fallout type attenuates its intensity linearly, scaled by the Dist value. This is the default setting, behaving the same as the default in previous Blender versions without Quad switched on, and consequently, as the new Lin/Quad Weighted option with Linear to **1.0** and Quad to **0.0** (TODO: Is it true? -- Mont29 13:42, 20 June 2009 (UTC)). This isnâ← t physically accurate, but can be easier to light with.

Constant



Render showing the Constant lamp falloff type effect with default settings.
This lamp fallout type does not attenuate its intensity with distance. This is useful for distant light sources like the sun or sky, which are so far away that their falloff isnâ← t noticeable. Sun and Hemi lamps always have constant falloff.

Sphere



Screenshot of the 3D view window, showing the Sphere light clipping circle.
The Sphere option restricts the light illumination range of a Lamp or Spot lamp, so that it will completely stop illuminating past an area once it reaches the number of Blender Units away from itself, as specified in the Dist field.

When the Sphere option is active, a dotted sphere will appear around the light source, indicating the demarcation point at which this light intensity will be null.

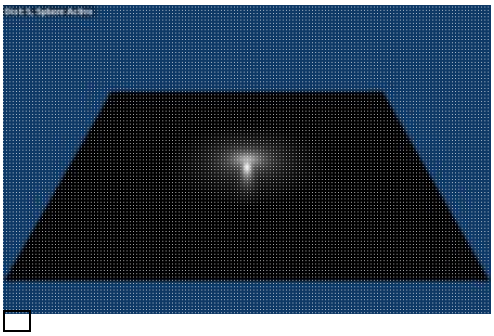
The Sphere option adds a term to the chosen attenuation law, whatever it is:

`Iâ← Â =Â IÂ Ã Â (DÂ -Â r)Â /Â DÂ` *`if`*`Â rÂ <Â D;Â 0Â` *`otherwise`*

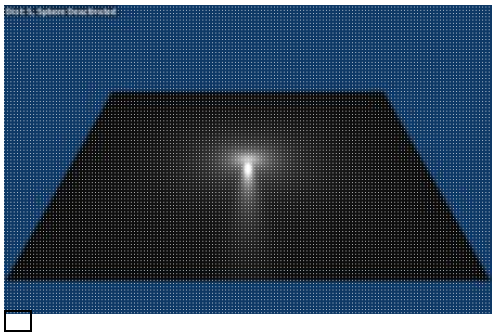
Where:

- `Iâ←` is the required Intensity of light (with the Sphere option activated).
- `I` is the intensity of light calculated by the chosen attenuation law (without the Sphere option).
- `D` is the current setting of the Dist field.
- `r` is the distance from the lamp where the light intensity gets measured.

See the graphic at the end of the description of the Lin/Quad Weighted attenuation optionâ←¦



Render showing the light attenuation of a Constant falloff light type with the Sphere option active.

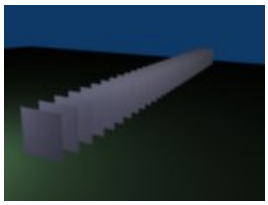


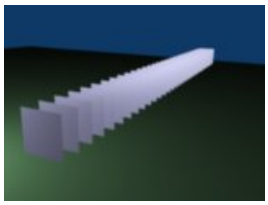
Render showing the light attenuation of a Constant falloff light type with the Sphere option deactivated.

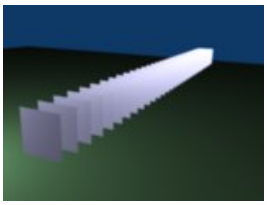
Examples

Distance

In this example, the Lamp has been set pretty close to the group of planes. This causes the light to affect the front, middle and rear planes more dramatically. Looking at (*Various Distance settings*), you can see that as the Dist is increased, more and more objects are becoming progressively brighter.


Dist: **10**.


Dist: **100**.


Dist: **1000**.

Various Distance settings (shadows disabled).

The Dist parameter is controlling where the light is falling â← Â at a linear rate by defaultÂ â← to half its original value from the lightâ← s origin. As you increase or decrease this value, you are changing where this half falloff occurs. You could think of Dist as the surface of a sphere and the surface is where the lightâ← s intensity has fallen to half its strength, in all directions. Note that the lightâ← s intensity continues to fall even after Dist. Dist just specifies the distance where half of the lightâ← s energy has weakened.

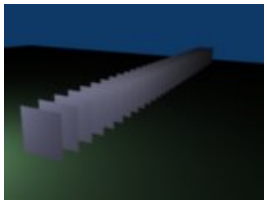
Notice in (*Dist: **1000***) that the farthest objects are very bright. This is because the falloff has been extended far into the distance which means the light is very strong when it hits the last few objects. It is not until **1000** units that the lightâ← s intensity has fallen to half of its original intensity.

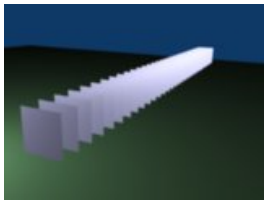
Contrast this with (*Dist: **10***), where the falloff occurs so soon that the farther objects are barely lit. The lightâ← s intensity has fallen by a half by time it even reaches the tenth object.

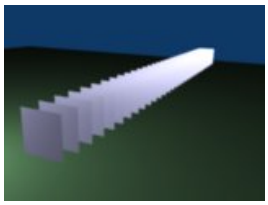
You may be wondering, why the first few planes appear to be dimmer? This is because the surface angle between the light and the objectâ← s surface normal are getting close to oblique. That is the nature of a Lamp light object. By moving the light infinitely far away you would begin to approach the characteristics of the Sun lamp type.

Inverse Square

Inverse Square makes the lightâ← s intensity falloff with a non-linear rate, or specifically, quadratic rate. The characteristic feature of using Inverse Square is that the lightâ← s intensity begins to fall off very slowly but then starts falling off very rapidly. We can see this in the (*Inverse Square selected*) images.


Inverse Square with **10**.


Inverse Square with **100**.


Inverse Square with **1000**.

Inverse Square selected (with the specified distances).

With Inverse Square selected, the Dist field is specifying where the light begins to fall faster, roughly speaking, see the light attenuation description for more info.

In (*Inverse Square with **10***), the lightâ← s intensity has fallen so quickly that the last few objects arenâ← t even lit.

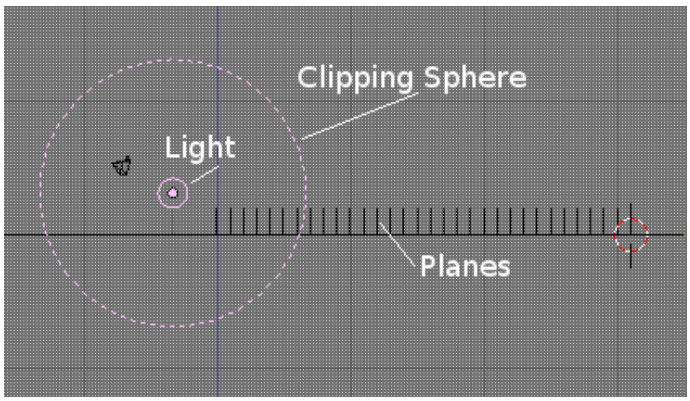
Both (*Inverse Square with **100***) and (*Inverse Square with **1000***) appear to be almost identical and that is because the Distance is set beyond the farthest objectâ← s distance which is at about **40 BU** out. Hence, all the objects get almost the full intensity of the light.

As above, the first few objects are dimmer than farther objects because they are very close to the light. Remember, the brightness of an objectâ← s surface is also based on the angle between the surface normal of

an object and the ray of light coming from the lamp.

This means there are at least two things that are controlling the surfaceâ ← s brightness: intensity and the angle between the light source and the surfaceâ ← s normal.

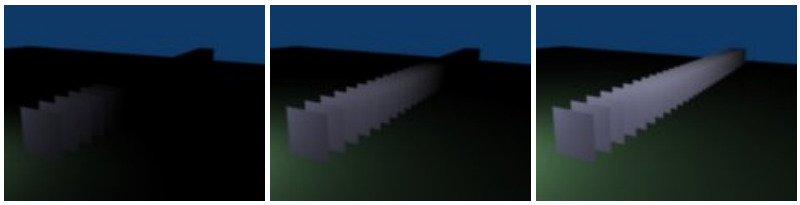
Sphere



Clipping Sphere.
Sphere indicates that the lightâ ← s intensity is null at the Dist distance and beyond, regardless of the chosen lightâ ← s falloff. In (*Clipping Sphere*) you can see a side view example of the setup with Sphere enabled and a distance of **10**.

Any objects beyond the sphere receive no light from the lamp.

The Dist field is now specifying both where the lightâ ← s rays become null, and the intensityâ ← s ratio falloff setting. Note that there is no abrupt transition at the sphere: the light attenuation is progressive (for more details, see the descriptions of the Sphere options and light attenuations above).



Sphere with **10**. Sphere with **20**. Sphere with **40**.
Sphere enabled with the specified distances, Inverse Linear light falloff.

In (*Sphere with **10***), the clipping sphereâ ← s radius is **10** units, which means the lightâ ← s intensity is also being controlled by **10** units of distance. With a linear attenuation, the lightâ ← s intensity has fallen very low even before it gets to the first object.

In (*Sphere with **20***), the clipping sphereâ ← s radius is now **20 BU** and some light is reaching the middle objects.

In (*Sphere with **40***), the clipping sphereâ ← s radius is now **40** units, which is beyond the last object. However, the light doesnâ ← t make it to the last few objects because the intensity has fallen to nearly **0**.

Hints

If a Lamp light is set to not cast shadows, it illuminates through walls and the like. If you want to achieve some nice effects like a fire, or a candle-lit room interior seen from outside a window, the Sphere option is a must. By carefully working on the Distance value you can make your warm firelight shed only within the room, while illuminating outside with a cool moonlight, the latter achieved with a Sun or Hemi light or both.

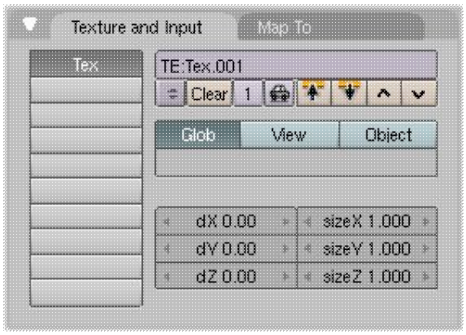
See Also

- Lamp light
- Spot light

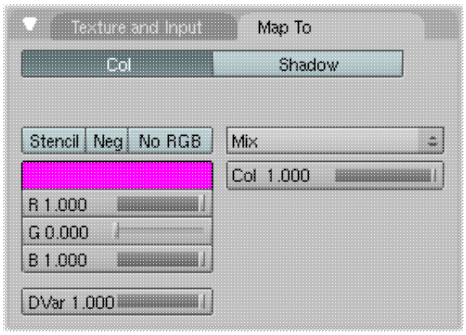
Lamps Textures

When a new lamp is added, it produces light in a uniform, flat colour. While this might be sufficient in simple renderings, more sophisticated effects can be accomplished through the use of textures. Subtle textures can add visual nuance to a lamp, while hard textures can be used to simulate more pronounced effects, such as a disco ball, dappled sunlight breaking through treetops, or even a projector. These textures are assigned to one of ten channels, and behave exactly like material textures, except that they affect a lampâ← s colour and intensity, rather than a materialâ← s surface characteristics.

Options



Texture and Input panel.



Map To panel.

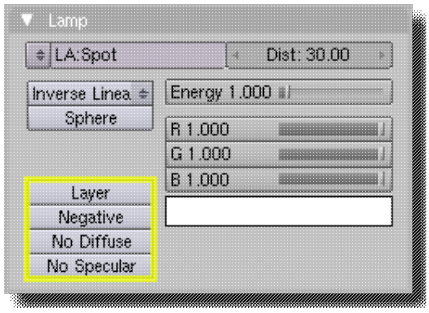
The lamp textures settings are regrouped in two panels. Here we will only talk about the few things that differ from usual material textures, see the Materials and Textures chapters for details about the standard options.

The two Texture and Map Input panels you are used to are here regrouped into the Texture and Input one. Except some GUI changes, the Texture part remains the same. However, youâ← ll note there are much less Map Input options â← Â you can only choose between Global, View or another Objectâ← s texture coordinates (as a lamp has no texture coordinates by itself), and you can scale or offset the texture.

The Map To panel is also a subset of its regular materialâ← s counterpart. You can only map a lamp texture to its regular, basic Color and/or to its Shadow color. As you can only affect colors, and a lamp has no texture coordinates on its own, the Nor/Var/Disp, and Warp options have disappeared.

What the Light Affects

Every lamp has a set of switches that control which objects receive its light, and how it interacts with materials.



Lamp panel with the light affecting options highlighted in yellow.

Layer
: Causes the lamp to only light objects on the same layer.

Negative
: The light produced by the lamp is **subtracted** from the one â← availableâ← on the surfaces it hits, which darkens these surfaces instead of brightening them.

No Diffuse
: Prevents the lamp from producing diffuse light (it doesnâ← t really â← lightâ← things).

No Specular
: Prevents the lamp from producing specular highlights.

Lamps Related Settings

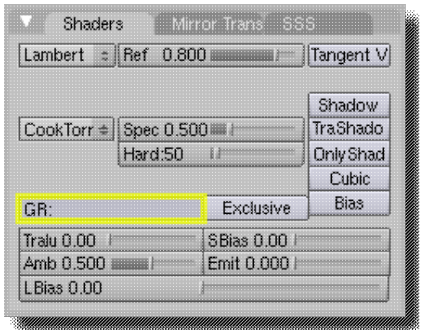
Mode: All modes

Panel: Shading context â Lamp sub-context (F5)

Here are some options closely related to light sources, without being lamps settings.

Lighting Groups

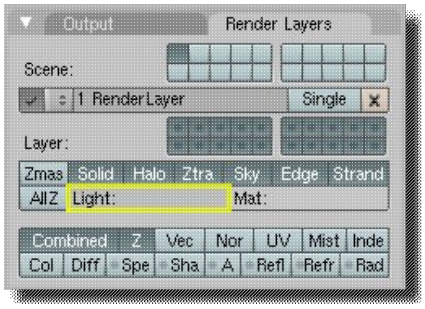
Materials



Limit lighting to lamps in this group highlighted in yellow.
By default, materials are lit by all lamps in all visible layers, but a material (and thus all objects using that material) can be limited to a single group of lamps. This sort of control can be incredibly useful, especially in scenes with complex lighting setups. To enable this, navigate to the Shading context, Material sub-context (F5), Shaders panel, and type the name of a lighting group in the GR input field. Note that the lighting group must be created before entering its name here!

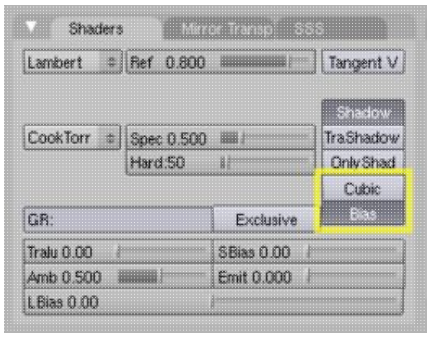
If the Exclusive button is enabled, lights in the specified group will *only* affect objects with this material.

Scene



Screenshot showing the Render Layers panel with the Light field highlighted in yellow.
Thereâ← s a similar control located in the Scene context (F10), Render sub-context, Render Layers panel. If a light group name is entered in this Light field, the scene will be lit exclusively by lamps in the specified group.

Other Options

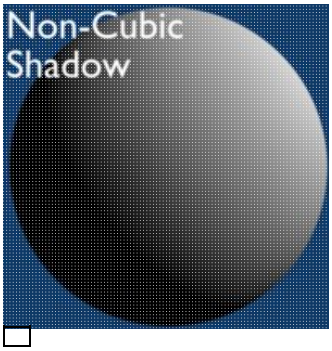


Cubic and Bias options in Shaders panel.

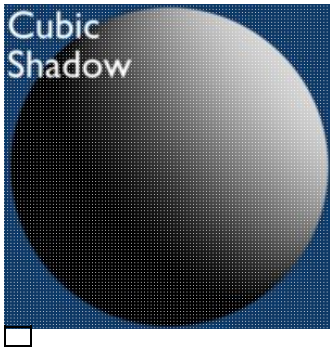
Back in the Material sub-context, if you look further down the side of the Shaders panel, you will see some other buttons related to lamps.

Cubic
: Uses cubic interpolation of diffuse values for smoother transitions between light and dark (see below *Without Cubic enabled* and *With Cubic enabled*).

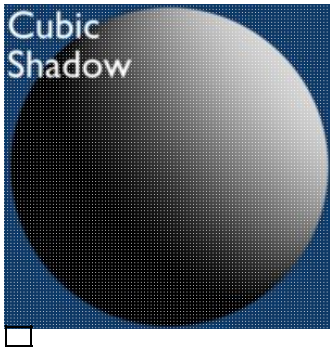
Bias
: Eliminates ray-traced shadow errors with the Phong specular shader.



Without Cubic enabled.



With Cubic enabled.



Animation switching between Non-Cubic and Cubic shadowing. Click to view animation.

See Also

- Lamps Introduction
- Shadows
- Materials Introduction

Shadows

Light wouldnâ← t even exist without its counterpart: shadows. Shadows are a darkening of a portion of an object because light is being partially or totally blocked from illuminating the object. They add contrast and volume to a scene; there is nearly no places in the real world without shadows, so to get realistic renders, you will need them. Blender supports the following kind of shadows:

1. Raytraced lamp shadows.
2. Buffered lamp shadows.
3. Ambient occlusion darkening.
4. Radiosity.

Ambient occlusion really isnâ← t a shadow based on light *per se*, but based on geometry. However, it does mimic an effect where light is prevented from fully and uniformly illuminating an object, so it is mentioned here. Also, it is important to mention ambient lighting, since increasing Ambient decreases the effect of a shadow. In that same vein, radiosity casts light from one object onto another, and if radiosity is high, the shadow will not appear either.

You can use a combination of ray-traced and buffer shadows to achieve different results. Even within raytraced shadows, different lamps cast different patterns and intensities of shadow. Depending on how you arrange your lamps, one lamp may wipe out or override the shadow cast by another lamp.

Shadows is one of those trifectas in Blender, where multiple things have to be set up in different areas to get results:

1. The lamp has to cast shadows (ability and direction).
2. Opaque object has to block light on its way (position and layer).
3. Another objectâ← s material has to receive shadows (Shadow and TraShadow enabled).
4. The render engine has to calculate shadows (Shadow for buffered shadows, Shadow and Ray for raytraced shadows).

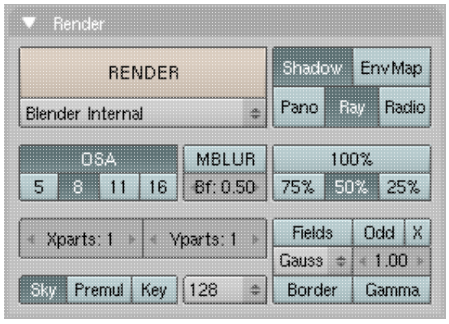
For example, the simple Lamp, Area, and Sun light has the ability to cast ray shadows, but not buffer shadows. The Spot light can cast both, whereas the Hemi light does not cast any. If a Sun lamp is pointing sideways, it will not cast a shadow from a sphere above a plane onto the plane, since the light is not traveling that way. All lamps able to cast shadows share some common options, described here.

Just to give you more shadow options (and further confuse the issue), lamps and materials can be set to respectively **only** cast and receive shadows, and not light the diffuse/specular aspects of the object. Also, render layers can turn on/off the shadow pass, and their output may or may not contain shadow informationâ←¦

Lamps: Raytraced Shadows

Mode: All modes

Panel: Shadow and Spot (Shading context, Lamp sub-context, F5)



Render panel.
Raytraced shadows produce very precise shadows with very low memory use, but at the cost of processing time. This type of shadowing is available to all lamp types except Hemi.

As opposed to buffered shadows, raytraced shadows are obtained by casting rays from a regular light source, uniformly and in all directions. The raytracer then records which pixel of the final image is hit by a ray light, and which is not. Those that are not are obviously obscured by a shadow.

Each light casts rays in a different way. For example, a Spot light casts rays uniformly in all directions within a cone. The Sun light casts rays from a infinitely distant point, with all rays parallel to the direction of the Sun light.

For each additional light added to the scene, with raytracing enabled, the rendering time increases. Raytraced shadows require more computation than buffered shadows but produce sharp shadow borders with very little memory resource usage.

To enable raytraced shadows, three actions are required:

- Enable shadows globally from the Scene context (F10), Render sub-context, using the Shadow button on the Render panel â← Â see (*Render panel*).
- Enable raytracing globally from the same panel, using the Ray button.
- Enable shadows for the light using the Ray Shadow button on the Shading context, Lamp sub-context, Shadow and Spot panel. This panel varies depending on the type of light.

All lamps able to cast raytraced shadows share some common options, described here.

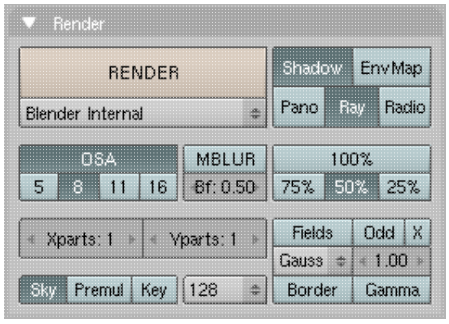
Raytraced shadows can be cast by the following types of lamp:

- Lamp lamp
- Spot lamp
- Area lamp
- Sun lamp

Lamps: Buffered Shadows

Mode: All modes

Panel: Shadow and Spot (Shading context, Lamp sub-context, F5)



Render panel.
Buffered shadows provides fast rendered shadows at the expense of precision and/or quality. Buffered shadows also require more memory resources as compared to raytracing. Using buffered shadows depends on your requirements. If you are rendering animations or canâ ← t wait hours to render a complex scene with soft shadows, buffer shadows are a good choice.

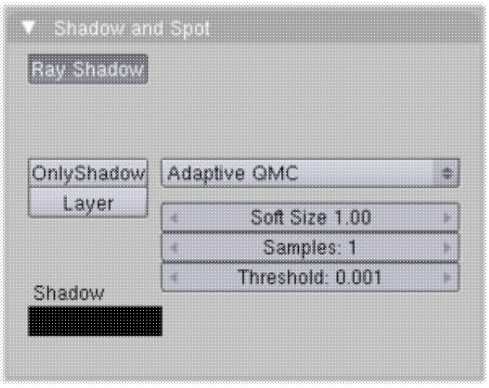
For a scanline renderer â ← Â and Blenderâ ← s built-in engine *is*, among other things, a scanline rendererÂ â ← , shadows can be computed using a *shadow buffer*. This implies that an â ← imageâ ← , as seen from the spot lampâ ← s point of view is â ← renderedâ ← and that the distance â ← Â in the imageÂ â ← for each point from the spot light is saved. Any point in the â ← renderedâ ← image that is farther away than any of those points in the spot lightâ ← s image is then considered to be in shadow. The shadow buffer stores this image data.

To enable buffered shadows two actions are required:

- Enable shadows globally from the Scene context (F10), Render sub-context, using the Shadow button on the Render panel â ← Â see (*Render panel*).
- Enable shadows for the light using the Buf Shadow button on the Shading context, Lamp sub-context (F5), Shadow and Spot panel. This panel varies depending on the type of light.

The only lamp able to cast buffered shadows is the Spot lamp.

Common Shadowing Lamps Options



Shadow and Spot panel.
All lamps able to cast shadows (Lamp, Spot, Area and Sun) share some options, described below:

Only Shadow

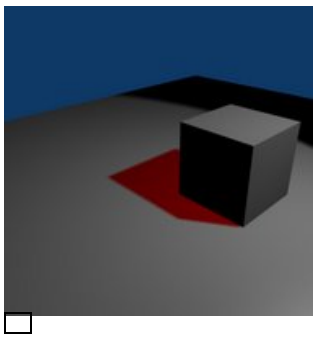
When the Only Shadow button is enabled, the light source will not illuminate an object but will generate the shadows that would normally appear.
This feature is often used to control how and where shadows fall by having a light which illuminates but has no shadow, combined with a second light which doesnâ← t illuminate but has Only Shadow enabled, allowing the user to control shadow placement by moving the â← Shadow Onlyâ← light around.

Layer

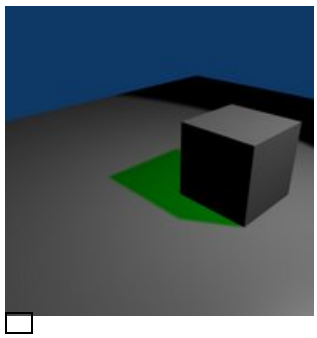
When this option is enabled, only the objects on the same layer as the light source will cast shadows.

Shadow color

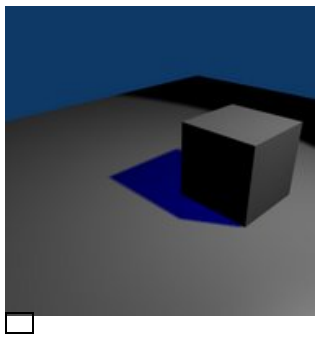
This color picker control allows you to choose the color of your cast shadows (black by default). The images below were all rendered with a white light and the shadow color was selected independently.



Red colored shadow example.



Green colored shadow example.



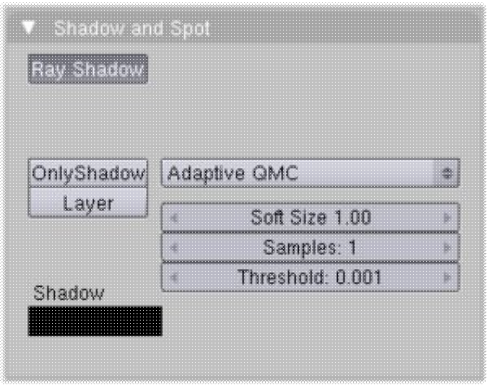
Blue colored shadow example.

Although you can select a pure white color for a shadow color, it appears to make a shadow disappear.

See Also

- Shadows
- Common Raytraced Options
- Lamp Light Raytraced Shadows
- Spot Light Raytraced Shadows
- Area Light Raytraced Shadows
- Sun Light Raytraced Shadows
- Spot Light Buffered Shadows

Lamps Raytraced Shadows



Shadow and Spot panel, raytraced shadows enabled, Adaptive QMC sample generator.
Most lamp types (Lamp, Spot and Sun) share the same options for the raytraced shadows generation, which are described below. Note that the Area lamp, even though using most of these options, have some specificities described in its own raytraced shadows page.

Ray Shadow

The Ray Shadow button enables the light source to generate raytraced shadows.
When the Ray Shadow button is selected, another set of options is made available, those options being:

Shadow sample generator type

This drop-down menu allows you to choose which algorithm is to be used to generate the samples that will serve to compute the raytraced shadows (for now, mainly two variants of Quasi-Monte Carlo, see below):

Constant QMC

The Constant QMC method is used to calculate shadow values in a very uniform, evenly distributed way. This method results in very good calculation of shadow value but it is not as fast as using the Adaptive QMC method, however Constant QMC is more accurate.

Adaptive QMC

The Adaptive QMC method is used to calculate shadow values in a slightly less uniform and distributed way. This method results in good calculation of shadow value but not as good as Constant QMC. The advantage of using Adaptive QMC is that it is in general much quicker while being not much worse than Constant QMC in terms of overall results.

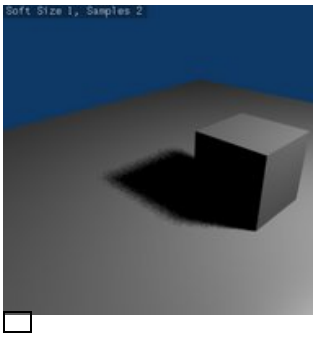
Samples

This slider sets the maximum number of samples (per pixel?) that both Constant QMC and Adaptive QMC will use to do their shadow calculations. The maximum value is **16** â← Â the real number of samples being actually the square of it, so setting a sample value of **3** really means 3^2Â =Â **9** samples will be taken.

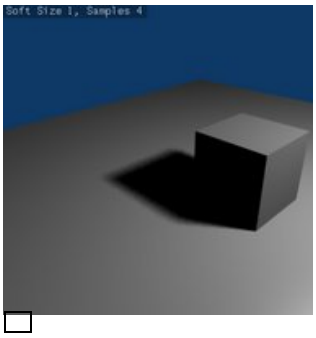
Soft Size

This slider determines the size of the fuzzy/diffuse/penumbra area around the edge of a shadow. Soft Size only determines the width of the soft shadow size, not how graded and smooth the shadow is. If you want a wide shadow which is also soft and finely graded, you must also set the number of samples in the Samples field higher than **1**, otherwise this field has no visible effect and the shadows generated will not have a soft edge. The maximum value for Soft Size is **100.0** (Blender Units, or percents?).
Below is a table of renders with different Soft Size and Samples settings showing the effect of various values on the softness of shadow edges:

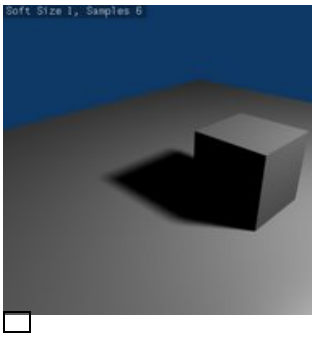
Lamps Raytraced Shadows



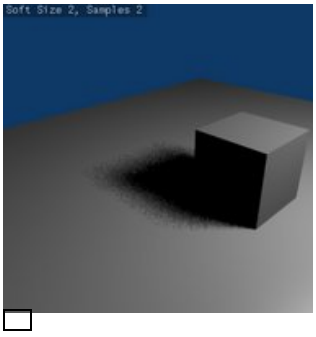
Soft Size: **1.0**, Samples: **2**.



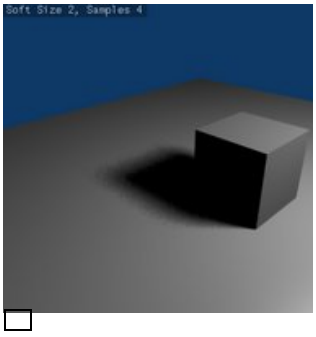
Soft Size: **1.0**, Samples: **4**.



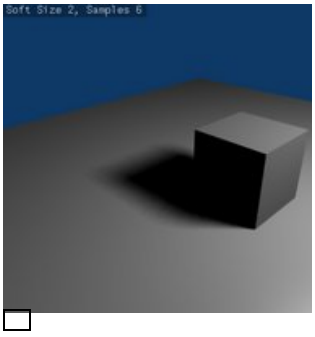
Soft Size: **1.0**, Samples: **6**.



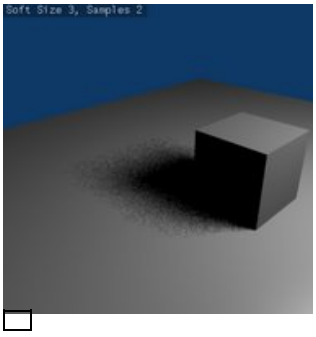
Soft Size: **2.0**, Samples: **2**.



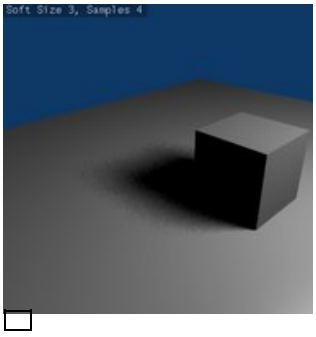
Soft Size: **2.0**, Samples: **4**.



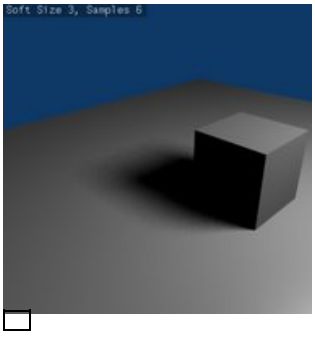
Soft Size: **2.0**, Samples: **6**.



Soft Size: **3.0**, Samples: **2**.

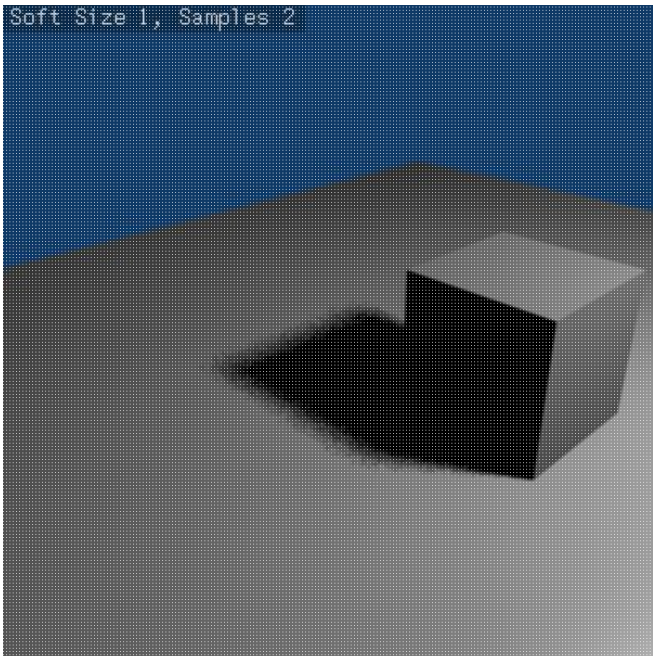


Soft Size: **3.0**, Samples: **4**.



Soft Size: **3.0**, Samples: **6**.

Below is an animated version of the above table of images showing the effects:



You may need to click on the image to see the animation.

Threshold

The Threshold field is used with the Adaptive QMC shadow calculation method. The value in the Threshold field is used to determine if Adaptive QMC shadow sample calculation can be skipped based on a threshold of how shadowed an area is already. The maximum Threshold value is **1.0**.

What is Quasi-Monte Carlo?

The Monte Carlo method is a method of taking a series of samples/readings of values (any kind of values, such as light values, color values, reflective states) in or around an area at random, so as to determine the correct actions to take in certain calculations which usually require multiple sample values to determine overall accuracy, of those calculations. The Monte Carlo methods tries to be as random as possible, this can often cause areas that are being sampled to have large irregular gaps in them (places that are not sampled/read), this in turn can cause problems for certain calculations (such as shadow calculation).

The solution to this was the Quasi-Monte Carlo method.

The Quasi-Monte Carlo method is also random, but tries to make sure that the samples/readings it takes are also better distributed (leaving less irregular gaps in its sample areas) and more evenly spread across an area. This has the advantage of sometimes leading to more accurate calculations based on samples/reading.

Volumetric Lighting

According to Wikipedia, volumetric lighting is as described below:

> Â«Â Volumetric lighting is a technique used in 3D computer graphics to add Tyndall-effect lighting to a rendered scene. The term seems to have been introduced from cinematography and is now widely applied to 3D modeling and rendering especially in the field of 3D gaming. It allows the viewer to see beams of light shining through the environment. Seeing sunbeams streaming through an open window is an example of volumetric lighting, also known as God rays.
> In volumetric lighting, the light cone emitted by a light source is modeled as a transparent object and considered as a container of a â← volumeâ← : as a result, light has the capability to give the effect of passing through an actual three dimensional medium (such as fog, dust, smoke, or steam) that is inside its volume, just like in the real world.Â Â»

A classic example is the search light with a visible halo/shaft of light being emitted from it as the search light sweeps around.

Blend file of spotlight animation.

By default Blender does not model this aspect of light. For example when Blender lights something with a Spot light, you see the objects and area on the floor lit but not the shaft/halo of light coming from the spotlight as it progresses to its target and would get scattered on the way.

The halo/shaft of light is caused in the real world by light being scattered by particles in the air, some of which get diverted into your eye and that you perceive as a halo/shaft of light. The scattering of light from a source can be simulated in Blender using various options, but by default is not activated.

The only lamp able to create volumetric effects is the Spot lamp (even thought you might consider some of the â← Sky & Atmosphereâ← effects of the Sun lamp as volumetric as wellâ←¦).

Lamps

Mode: All modes

Panel: Preview (Shading context, Lamp sub-context, F5)

Hotkey: â § ShiftA to add new

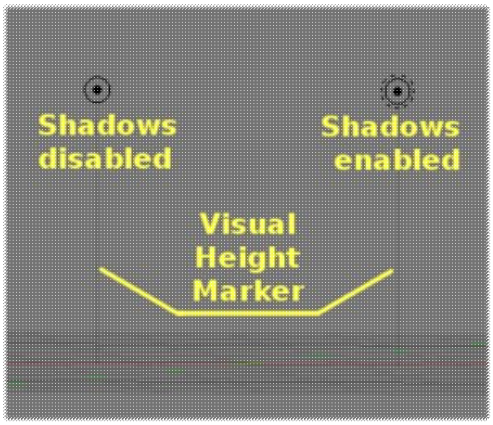
Menu: Add Â» Lamp

Blender comes equipped with five different lamp types, each with its own unique strengths and limitations. Here are the available lamps:

- Lamp is an omnidirectional point light source, similar to a light bulb.
- Spot is a directional point light source, similar toâ ←¦ a spot!
- Area is a source simulating area producing light, as windows, neons, TV screensâ ←¦
- Hemi simulates a very wide and far away light source, like the sky.
- Sun simulates a very far away and punctual light source, like the sun.

You can add new lamps to a scene using the Add menu in the top header, or by using the add menu in the toolbox (Space Â» Add Â» Lamp), or â § ShiftA.

Once added, a lampâ ← s position is indicated in the 3D View by a solid dot in a circle, but most types also feature dashed wire-frames that help describe their orientation and properties. While each type is represented differently, there are some visual indicators common to all of them:



A default Lamp, showing the visual height and shadow markers.

Shadows
: If shadows are enabled, an additional dashed circle is drawn around the solid circle. This makes it easier to quickly determine if a lamp has shadows enabled.

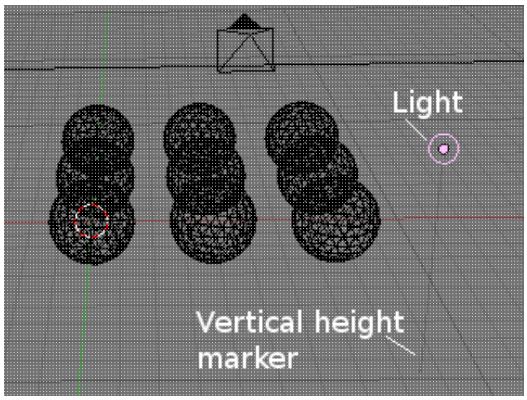
Vertical Height Marker
: This is a dim grey line, which helps locate the lampâ ← s position relative to the global X-Y plane. The lineâ ← s transparency can be adjusted in the Theme section of the User Preferences window, with the Alpha value of the 3D ViewÂ â Lamp item (if you use the default theme, you must make a copy of it to be able to modify this setting!).

Lamp

Mode: All modes

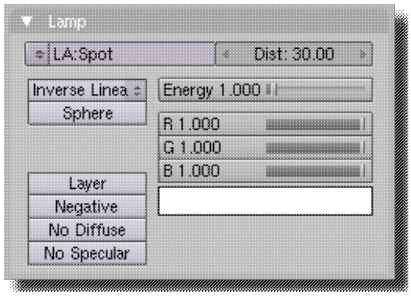
Panel: Lamp (Shading context, Lamp sub-context, F5)

Description



Lamp light.
The Lamp lamp is an omnidirectional point of light, that is, a point radiating the same amount of light in all directions. Itâ← s visualized by a plain, circled dot, (*Lamp light*). Being a point light source, the direction of the light hitting an objectâ← s surface is determined by the line joining the lamp and the point on the surface of the object itself. Light intensity/energy decays based on (among other variables) distance from the Lamp to the object. In other words, surfaces that are further away are darker.

Options



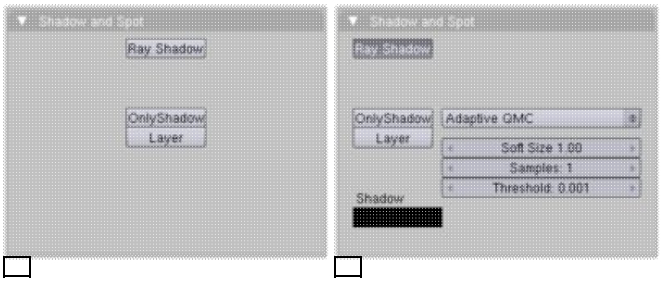
Lamp panel.

Common options
: The Dist, Energy and Color settings are common to most types of lamps, and are described here.
The Layer, Negative, No Diffuse and No Specular settings control what the lamp affect, as described in this page.

Light attenuation drop-down list and Sphere button
: These settings control how the light of the Lamp decays with the distance. See this page for details.

â← Shadow and Spotâ← Panel

When a Lamp light source is selected the Shadow and Spot panel has the following default layout:

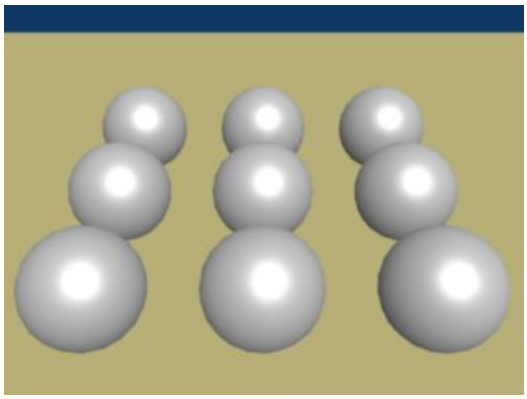


Without ray shadows. With ray shadows, Adaptive QMC sample generator.

The Shadow and Spot panel when Lamp light source is selected.

As its only purpose for this lamp is about raytraced shadows, it is described in details here.

Examples



Render example.

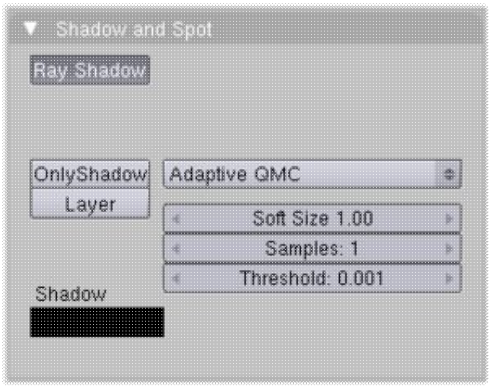
Notice in (*Render example*) that the light is gradually diminishing for each sphere that is farther away from the light source as compared to the Sun render example, where the lightâ← s intensity is constant (i.e. never fades with distance).

You can find more useful and complete examples about lamps in the â← lighting rigsâ← page.

See Also

- Light attenuation
- Shadows
- Lamp raytraced shadows
- Lighting rigs

Lamp Raytraced Shadows



Shadow and Spot panel for Lamp light source, raytraced shadows enabled, Adaptive QMC sample generator.
The Lamp light source can only cast raytraced shadows. It shares with other lamp types common shadows options described here.

The raytraced shadows settings of this lamp are shared with other ones, and are described here.

See Also

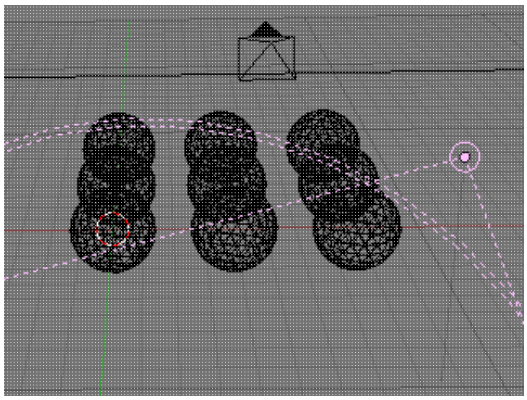
- Shadows
- Common lamp shadows options
- Common lamp ray shadows options
- Lamp light source

Spot

Mode: All modes

Panel: Lamp (Shading context, Lamp sub-context, F5)

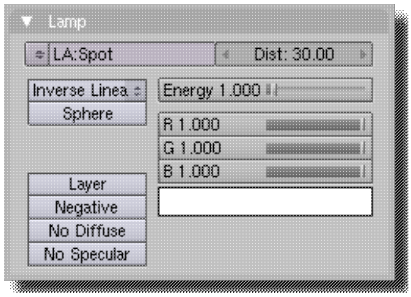
Description



Spot lamp.
A Spot lamp emits a cone shaped beam of light from the tip of the cone, in a given direction.

The Spot light is the most complex of the light objects and indeed, for a long time, among the most used thanks to the fact that it was the only one able to cast shadows. Today, with the integration of a ray tracer within the internal render engine of Blender, all lamps can cast shadows (except Hemi). Even so, Spot lampsâ← shadow buffers are much faster to render than raytraced shadows, especially when blurred/softened, and spot lamps also provide other functionality such as â← volumetricâ← halos.

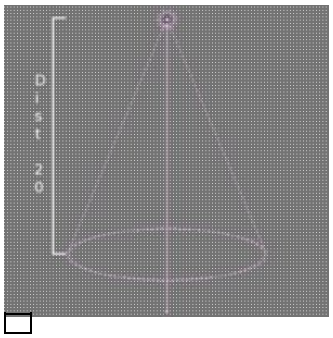
Options

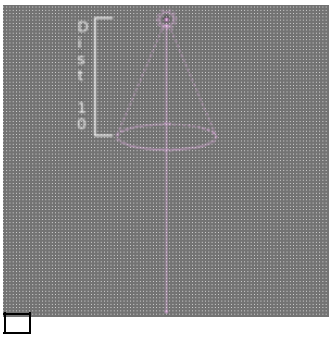


Lamp panel.

Common options

The Dist, Energy and Color settings are common to most types of lamps, and are described here.
The Layer, Negative, No Diffuse and No Specular settings control what the lamp affect, as described in this page.





Spotlight with Dist value of '*20.0*. Spotlight with Dist value of **10.0**.

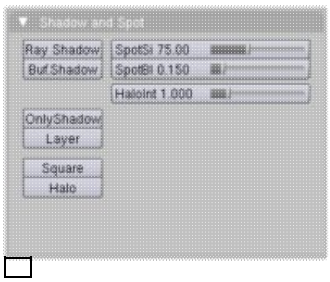
Changing the Dist field value when using a Spot light also changes the appearance of the spotlight as displayed in the 3D Viewport:

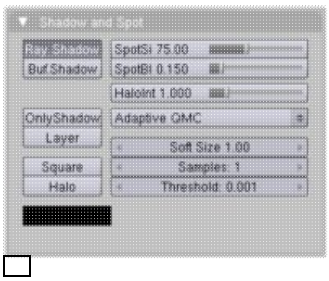
Light attenuation drop-down list and Sphere button

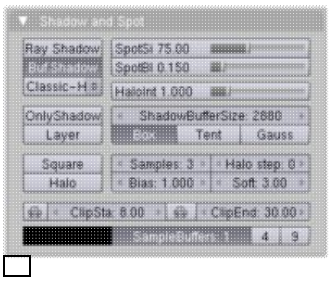
These settings control how the light of the Spot decays with the distance. See this page for details.

â← Shadow and Spotâ← Panel

When a Spot lamp is selected the following default layout for the Shadow and Spot panel is shown:



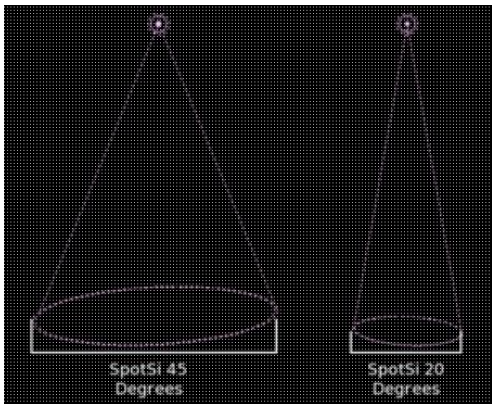




Without shadow. With raytraced shadows. With buffered shadows.

The Shadow and Spot panel when Spolt lamp lighting is selected.

Options



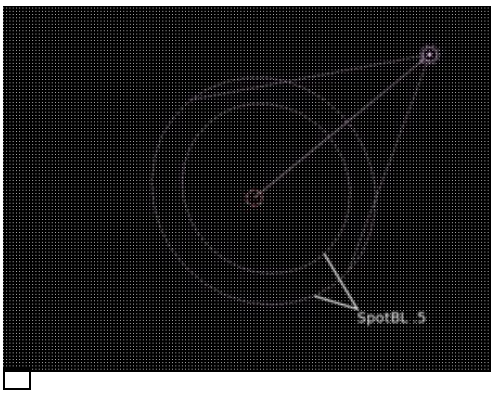
SpotSi settings at **45.0** and **20.0** degrees.

SpotSi

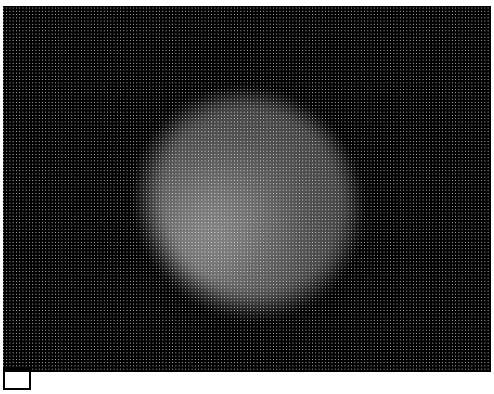
The SpotSi (for â← SpotSizeâ←) slider controls the size of the outer cone of a Spot, which largely controls the circular area a Spot light covers. This slider controls in fact the angle at the top of the lighting cone, and can be between **1.0Â°** et **180.0Â°**.

SpotBL

SpotBL (for â← SpotBLurâ←) slider controls the inner cone of the Spot. The SpotBL value can be between **0.0** and **1.0**. The value is proportional and represents that amount of space that the inner cone should occupy inside the outer cone (SpotSi).

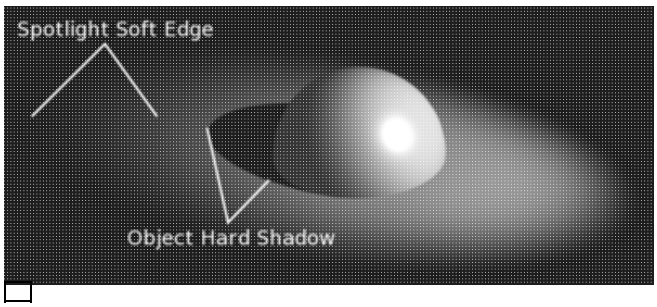


3D View of Spot lamp with a SpotBL setting of **0.5**.



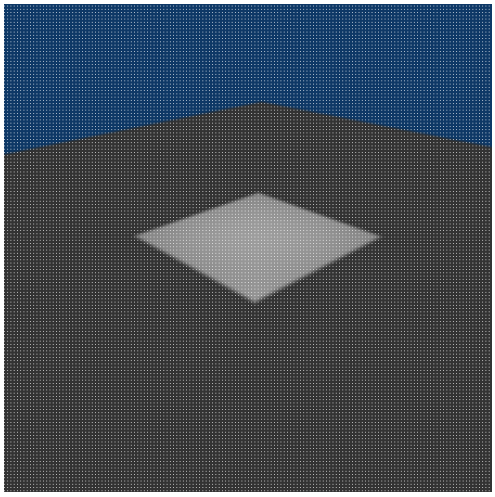
Render of Spot lamp with a SpotBL setting of **0.5**.

The inner cone boundary line indicates the point at which light from the Spot will start to blur/soften, before this point its light will mostly be full strength. The larger the value of the SpotBL the more blurred/soft the edges of the spotlight will be, and the smaller the inner cones circular area will be (as it starts to blur/soften earlier).
To make the Spot have a sharper falloff rate and therefore less blurred/soft edges, decrease the value of SpotBL. Setting SpotBL to **0.0** results in very sharp spotlight edges, without any transition between light and shadow.
The falloff rate of the Spot lamp light is a ratio between the SpotBL and SpotSi values, the larger the circular gap between the two, the more gradual the light fades between SpotBL and SpotSi.
SpotBL and SpotSi only control the Spot light coneâ← s aperture and softness (â← radialâ← falloff), they do not control the shadowâ← s softness as shown below.



Render showing the soft edge spotlighted area and the sharp/hard object shadow.

Notice in the picture above that the objectâ← s shadow is sharp as a result of the raytracing, whereas the spotlight edges are soft. If you want other items to cast soft shadows within the Spot area, you will need to alter other shadow settings.



Spotlight with a Square light area.

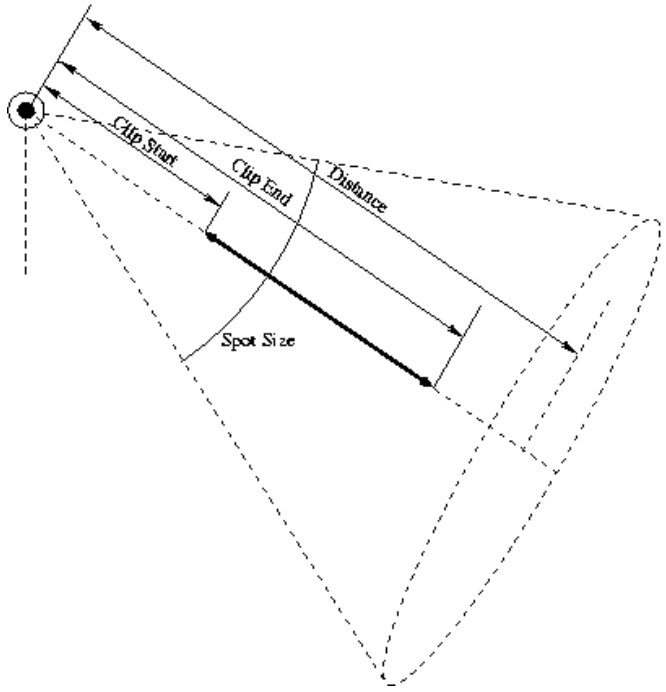
Square
: The Square button makes a Spot light cast a square light area, rather than the default circular one.

Shadows and Halos

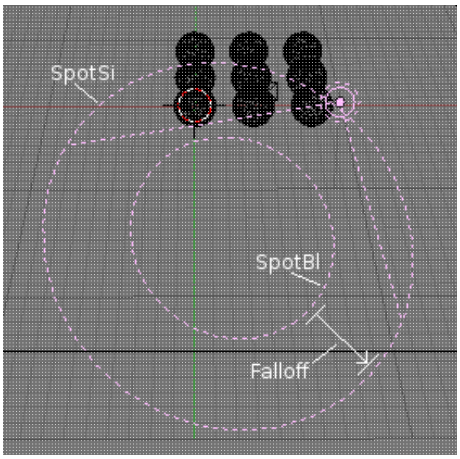
The other options in this panel are about shadows and volumetric effects. The Spot light source supports raytraced and buffered shadows, along with halos, follow these links for in-depth descriptions of these topics.

Technical Details

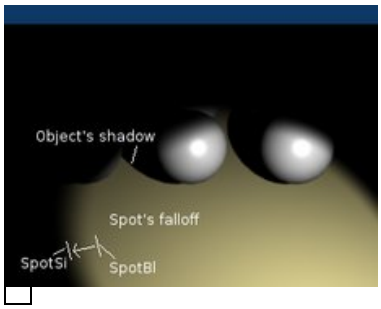
(*Spot light scheme*) shows the relationship between the lightâ← s properties and how they relate physically.



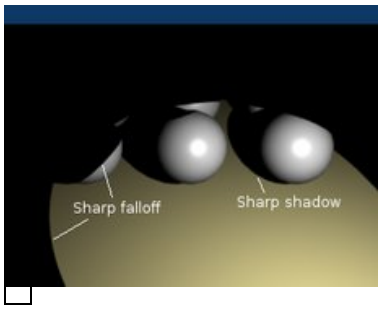
Spot lamp scheme.



SpotSi and SpotBl.

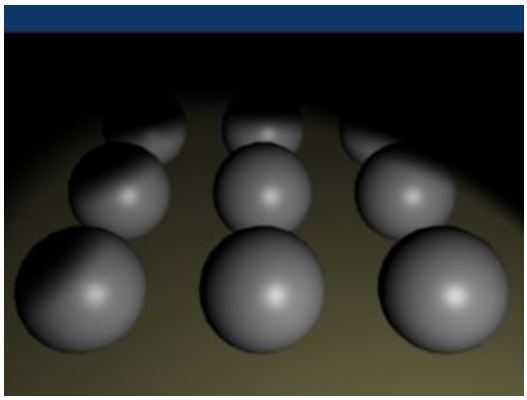


Shadow comparison.



Sharp falloff.

Examples

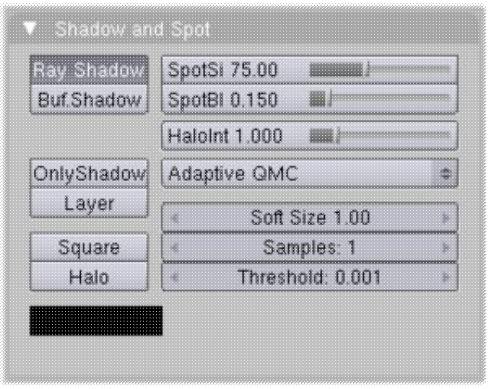


Spot lamp render example.
You can find more useful and complete examples about lamps in the â← lighting rigsâ← page.

See Also

- Light attenuation
- Shadows
- Spot raytraced shadows
- Spot buffered shadows
- Spot halos
- Lighting rigs

Spot Raytraced Shadows



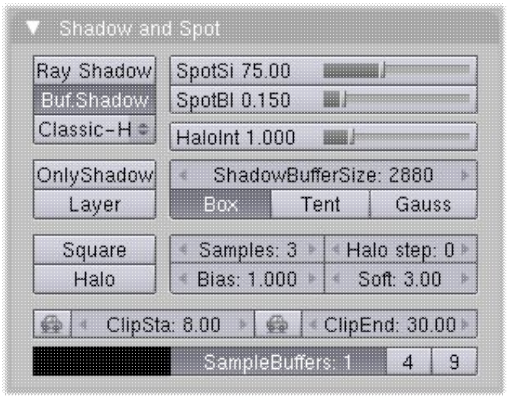
The Shadow and Spot panel for a Spot lamp, with raytraced shadows activated.
Spotlights can use either raytraced shadows or buffered shadows. Either of the two can provide various extra options. Raytraced shadows are generally more accurate, with extra capabilities such as transparent shadows, although they are quite slower to render. Buffered shadows are more complex to set up and involve more faking, but the speed of rendering is a definite advantage. Anyway, it shares with other lamp types common shadows options described here.

The raytraced shadows settings of this lamp are shared with other ones, and are described here.

See Also

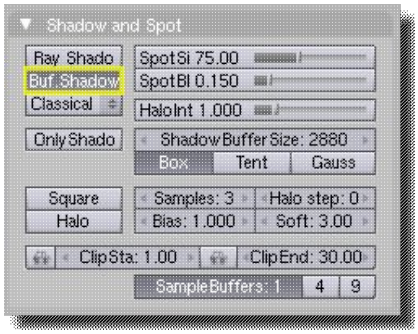
- Shadows
- Common lamp shadows options
- Common lamp ray shadows options
- Spot light source
- Spot buffered shadows
- Spot halos

Spot Buffered Shadows



The Shadow and Spot panel for a Spot lamp, when using buffered shadows (Classical or Classic-Halfway shadow generation method).
Spotlights can use either raytraced shadows or buffered shadows. Either of the two can provide various extra options. Raytraced shadows are generally more accurate, with extra capabilities such as transparent shadows, although they are quite slower to render. Buffered shadows are more complex to set up and involve more faking, but the speed of rendering is a definite advantage. Anyway, it shares with other lamp types common shadows options described here.

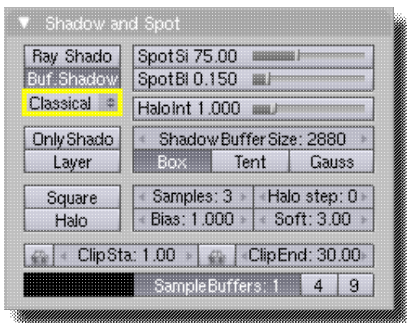
Buffered Shadows



Shadow and Spot panel Layout with Buf. Shadow button highlighted in yellow.

Buf. Shadow
: When the Buf. Shadow button is activated, the currently selected Spot light generates shadows, using a â← shadow bufferâ← rather than using raytracing.

When the Buf. Shadow button is activated, various extra options and buttons appear in the Shadow and Spot panel. A description of most of these options is given below:



Shadow buffer generation type field highlighted in yellow.

Shadow buffer generation type

Blender has more than one way to generate buffered shadows. The shadow buffer generation type drop-down selector controls which generator to use.

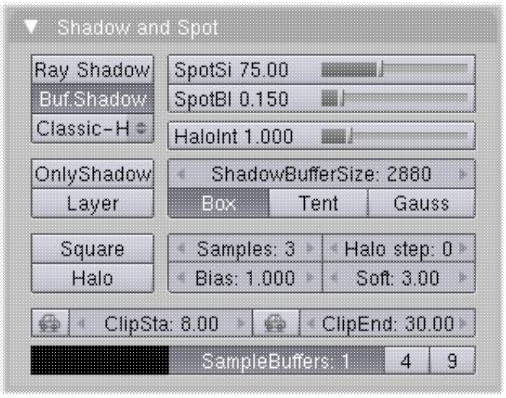
There are three shadow generation types, those being:

◊ Classical
◊ Classic-Halfway
◊ Irregular

For more information on the different shadow generation methods see these links:

- www.blender.org/development/release-logs/blender-243/irregular-shadow-buffer/
- www.blendernation.com/2006/10/15/blender-gets-irregular-shadow-buffers/
- www.blender.org/development/release-logs/blender-243/shadow-buffer-halfway-average/

â← Classicalâ← and â← Classic-Halfwayâ← generation methods



Shadow buffer options when using the Classical or Classic-Halfway shadow generation method.
Classical shadow generation used to be the Blender default and unique method for generation of buffered shadows. It used an older way of generating buffered shadows, but it could have some problems with accuracy of the generated shadows and can be very sensitive to ShadowBufferSize, different Bias values, and all the self-shadowing issues that brings up. It appears that the Classical method of generating shadows is in the position of being obsoleted and is really only still present to works with older versions of Blender, Classic-Halfway should probably be used instead.

Classic-Halfway is an improved shadow buffering method and is currently the default option selected in Blender. It works by taking an averaged reading of the first and second nearest Z depth values allowing the Bias value to be lowered and yet not suffer as much from self-shadowing issues. Not having to increase Bias values helps with shadow accuracy, because large Bias values can mean small faces can lose their shadows, as well as preventing shadows being overly offset from the larger Bias value.

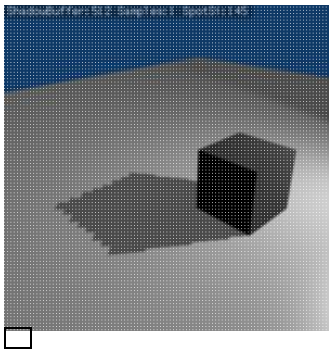
Classic-Halfway doesnâ← t work very well when faces overlap, and biasing problems can happen.

Currently the Halo Step option doesnâ← t work well in some cases. Especially when using planes (not volumes), errors can be introduced.

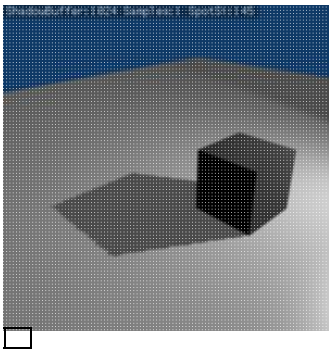
Here are now the options specific to these generation methods:

ShadowBufferSize

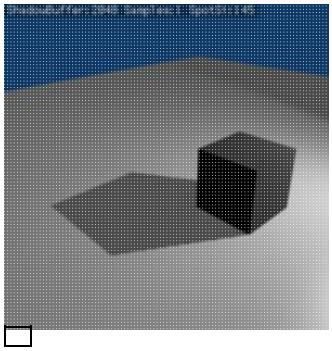
The ShadowBufferSize numeric field can have a value from **512** to **10240**. ShadowBufferSize represents the resolution used to create a shadow map. This shadow map is then used to determine where shadows lay within a scene.
As an example, if you have a ShadowBufferSize with a value of **1024**, you are indicating that the shadow data will be written to a buffer which will have a square resolution of **1024Ã 1024** pixels/samples from the selected spotlight.
The higher the value of ShadowBufferSize, the higher resolution and accuracy of the resultant shadows, assuming all other properties of the light and scene are the same, although more memory and processing time would be used. The reverse is also true â← Â if the ShadowBufferSize value is lowered, the resultant shadows can be of lower quality, but would use less memory and take less processing time to calculateâ←¦



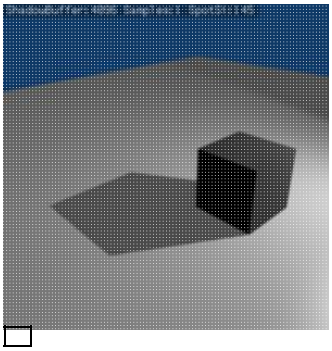
ShadowBufferSize: **512**.



ShadowBufferSize: **1024**.



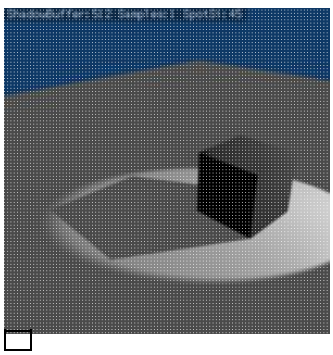
ShadowBufferSize: **2048**.

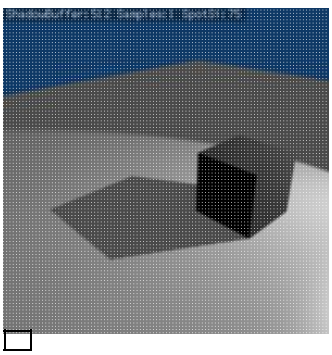


ShadowBufferSize: **4096**.

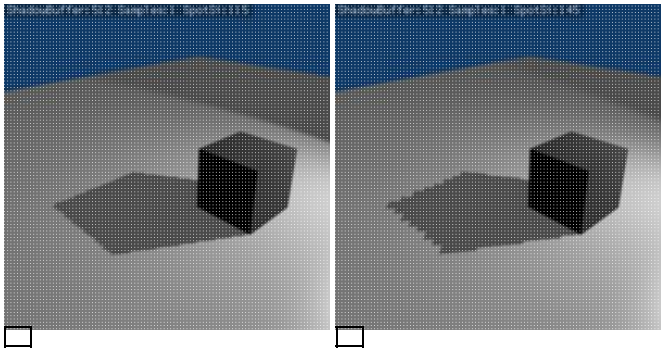
Buffered shadow examples for various ShadowBufferSize values (SpotSi to **145.0Â°**, no filtering â← Â Sample set to **1**Â â← , to make it easier to see the quality degradation of the shadows).

As well as the ShadowBufferSize value affecting the quality of generated shadows, another property of Spot lamps that affects the quality of their buffered shadows is the size of the spotlights lighted area (the SpotSi value), as you can see below:





SpotSi: **45.0Â°**. SpotSi: **75.0Â°**.



SpotSi: **115.0Â°**. SpotSi: **145.0Â°**.

Buffered shadow examples for various SpotSi values, with ShadowBufferSize set to **512**, and no filtering (Sample set to **1**).

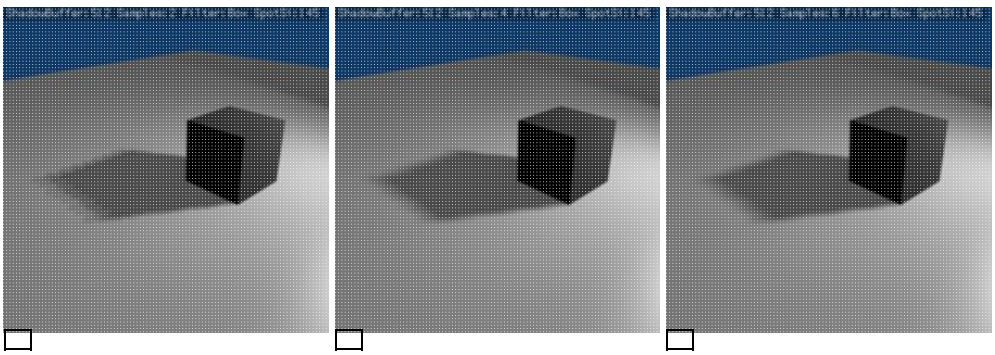
As the SpotSi value is increased, you can see that the quality of the cast shadows degrades.
This happens because when the Spot lighted area is made larger (by increasing SpotSi), the shadow buffer area has to be stretched and scaled to fit the size of the new lighted area. The ShadowBufferSize resolution is not altered to compensate for the change in size of the spotlight, so the quality of the shadows degrades. If you want to keep the generated shadows the same quality, as you increase the SpotSi value, you also need to increase the ShaodwBufferSize value.
The above basically boils down to:
If you have a spotlight that is large you will need to have a larger ShadowBufferSize to keep the shadows good quality. The reverse is true also â← Â the quality of the generated shadows will usually improve (up to a point) as the Spot lamp covers a smaller area.

Filtering algorithms

The three mutually-exclusive Box, Tent, and Gauss buttons control what filtering algorithm to use to anti-alias the buffered shadows.
They are closely related to the Samples numeric field, as when this setting is set to **1**, shadow filtering is disabled, so none of these buttons will have any effectâ←¦

Box

The buffered shadows will be anti-aliased using the â← boxâ← filtering method.
This is the original filter used in Blender. It is relatively low quality and is used for low resolution renders, as it produces very sharp anti-aliasing. When this filter is used, it only takes into account oversampling data which falls within a single pixel, and doesnâ← t take into account surrounding pixel samples. It is often useful for images which have sharply angled elements and horizontal/vertical lines (according to arkavision.com/?page_id=125).

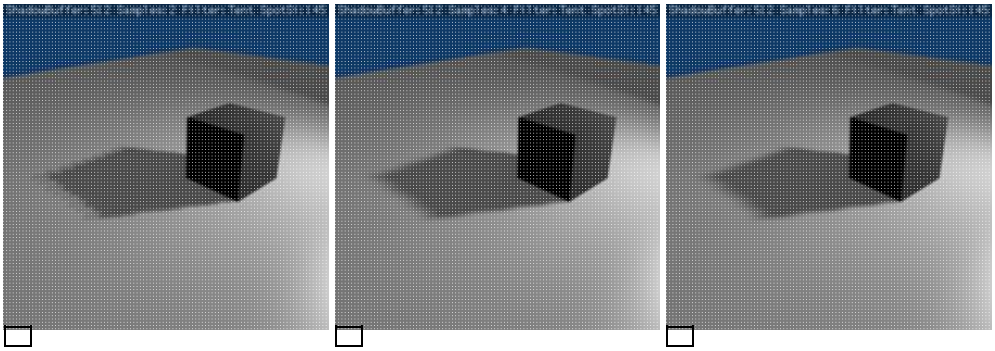


Samples: **2**. Samples: **4**. Samples: **6**.

Buffered shadow examples for various Samples values, with the Box filtering method (SpotSi: **145.0Â°**, ShadowBufferSize: **512**).

Tent

The buffered shadows will be anti-aliased using the â← tentâ← filtering method.
It is a simple filter that gives sharp results, an excellent general purpose filtering method. This filter also takes into account the sample values of neighboring pixels when calculating its final filtering value.

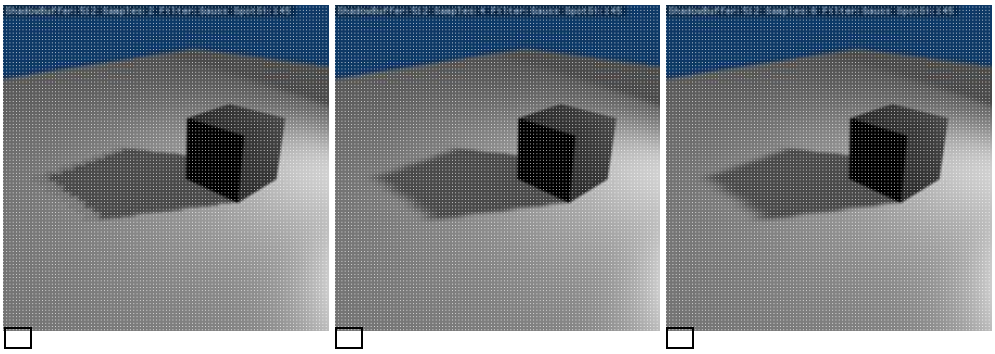


Samples: **2**. Samples: **4**. Samples: **6**.

Buffered shadow examples for various Samples values, with the Tent filtering method (SpotSi: **145.0Â°**, ShadowBufferSize: **512**).

Gauss

The buffered shadows will be anti-aliased using the â← gaussianâ← filtering method.
It produces a very soft/blurry anti-aliasing. As result, this filter is excellent with high resolution renders.



Samples: **2**. Samples: **4**. Samples: **6**.

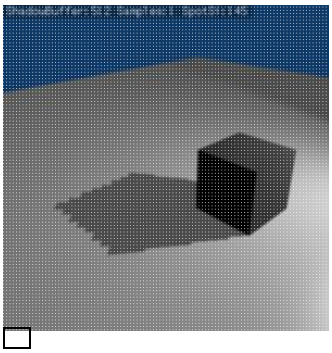
Buffered shadow examples for various Samples values, with the Gauss filtering method (SpotSi: **145.0Â°**, ShadowBufferSize: **512**).

The following links will give more information on the various filtering/distribution methods and their uses:

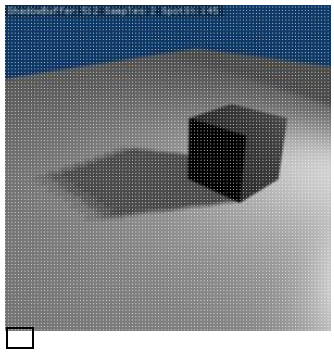
◊ The oversampling page in Render chapter
◊ arkavision.com/?page_id=125

Samples

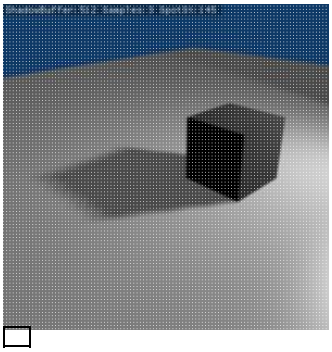
The Samples numeric field can have a value between **1** and **16**. It controls the number of samples taken per pixel when calculating shadow maps.
The higher this value, the more filtered, smoothed and anti-aliased the shadows cast by the current lamp will be, but the longer they will take to calculate and the more memory they will use. The anti-aliasing method used is determined by having one of the Box, Tent or Gauss buttons activated (see above).
Having a Samples value of **1** is similar to turning off anti-aliasing for buffered shadows.



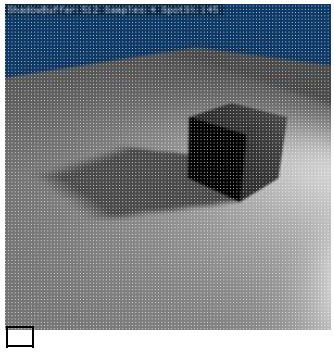
Samples: **1** (i.e. no anti-aliasing).



Samples: **2**.



Samples: **3**.

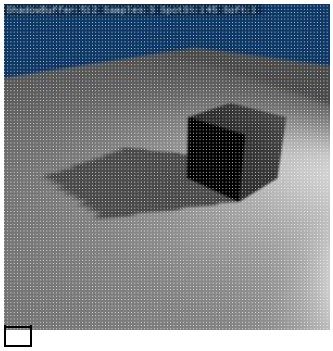


Samples: **4**.

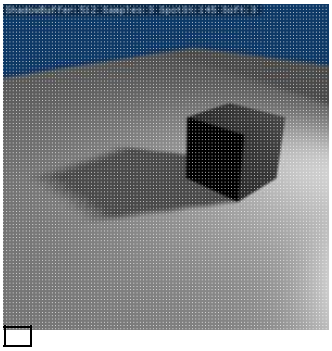
Buffered shadow examples for various Samples values (Box filtering method?, SpotSi: **145.0Â°**, ShadowBufferSize: **512**).

Soft

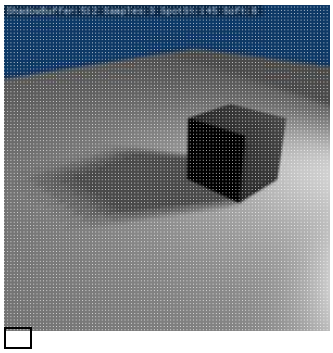
The Soft numeric field can have a value between **1.0** and **100.0**. It indicates how wide an area is sampled when doing anti-aliasing on buffered shadows. The larger the Soft value, the more graduated/soft the area that is anti-aliased/softened on the edge of generated shadows.



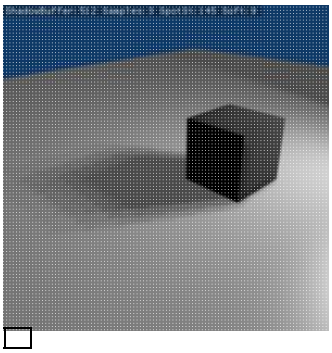
Soft: **1.0**.



Soft: **3.0**.



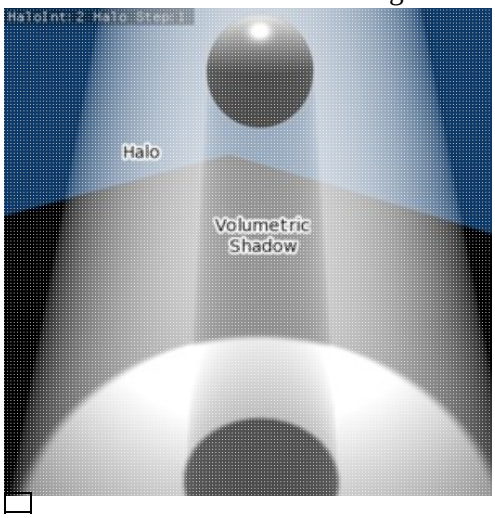
Soft: **6.0**.



Soft: **9.0**.

Buffered shadow examples for various Soft values (Box filtering method?, Samples: **3**, SpotSi: **145.0Â°**,

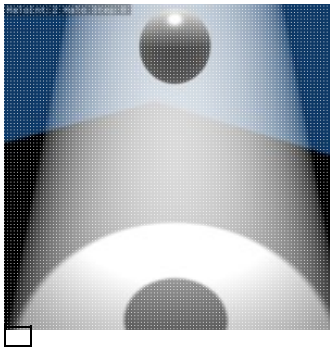
ShadowBufferSize: **512**).
The SpotSi and ShadowBufferSize have been chosen to exaggerate the blocky aliasing, and emphasize the Soft settingâ← s effectsâ←¦



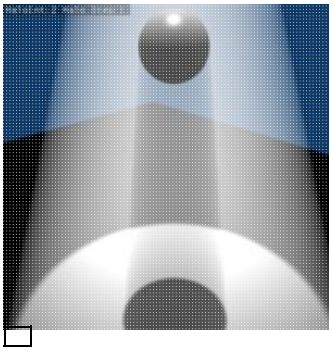
Layout of halo and volumetric shadow.

Halo Step

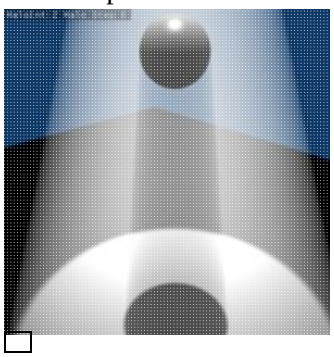
Halo Step can have a value between **0** and **12**. It is used to determine weather this Spot will cast volumetric shadows, and what quality those volumetric shadows will be.
(*Layout of halo and volumetric shadow*) is an example render which shows a volumetric shadow being cast.
For volumetric shadows to work, you must have the Halo button activated, and a high enough HaloInt value, so that the cast volumetric shadow is visible. Once these conditions have been met, the Halo Step value can be altered to change the quality of the volumetric shadows.
If Halo Step is set to a value of **0**, then no volumetric shadow will be generated.
Unlike most other controls, as the Halo Step value increases, the quality of volumetric shadows decreases (but takes less time to render), and *vice-versa*.



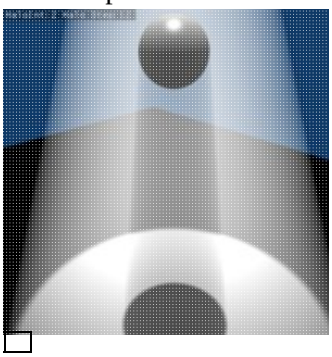
Halo Step: **0**.



Halo Step: **1**.



Halo Step: **6**.



Halo Step: **12**.

Volumetric shadow examples for various Halo Step values (HaloInt: **2.0**).

HaloStep values:

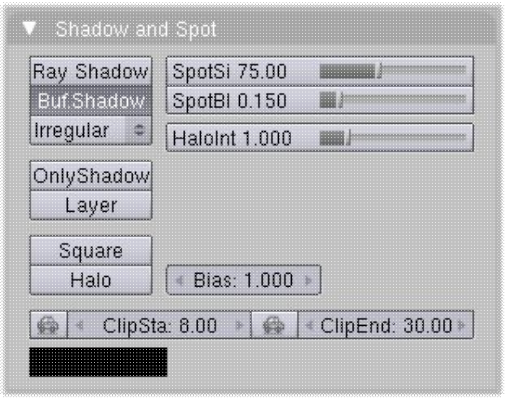
A value of **8** for Halo Step is usually a good compromise between speed and accuracy.

SampleBuffers: 1, 4, 9

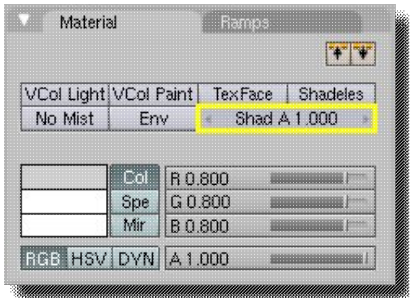
The SampleBuffers setting can be set to values **1**, **4** or **9**, and represents the number of shadow buffers that will be used when doing anti-aliasing on buffered shadows.
This option is used in special cases, like very small objects which move and need to generate really small shadows (such as strands). It appears that normally, pixel width shadows donâ← t anti-alias properly, and that increasing ShadowBufferSize doesnâ← t help much.
So this option allows you to have a sort of extra sample pass, done above the regular one (the one controlled by the Box/Tent/Gauss, Samples and Soft settings).
The **1** value will disable this option.
Higher values will produce a smoother anti-aliasing â← Â but be careful: using a SampleBuffers of **4** will require four times as much memory and process time, and so on, as Blender will have to compute that number of sample buffersâ← ¦
The â← pixelsâ← of each of the sample buffers are slightly offset, in a similar way as the standard rendering OSA (TODO: or isnâ← t it rather like the MBLUR temporal â← bluringâ← Â ?).
TODO:

- ◊ Find out how the sample buffers are really generated (spacial offsetting, temporal offsetting, or both?).
- ◊ Find out the method/algorithm used to blend/mix/combine the computed shadow buffers back into a single one.
- ◊ Explain why this method is better, for small shadows, than the standard, single-buffer one.

â← Irregularâ← generation method



Shadow buffer options when using the Irregular shadow generation method.
Irregular shadow method is used to generate sharp/hard shadows that are placed as accurately as raytraced shadows. This method offers very good performance because it can be done as a multi-threaded process.



The Shad A numeric field highlighted in yellow.
This method supports transparent shadows. To do so, first select the object which will receive the transparent shadow. Then alter the Shad A value in its Material panel (Shading context, Material sub-context, F5). This will only work when Irregular shadow buffer lighting is used.

TODO:

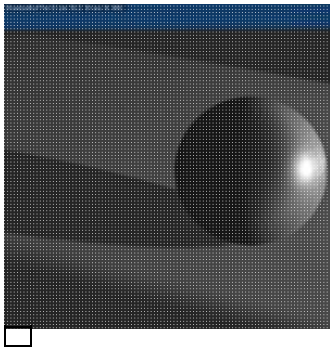
- Find how this method is working (and why it has nearly no extra option).

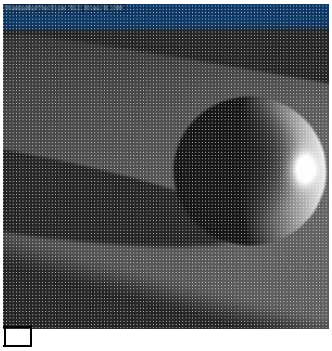
Common options

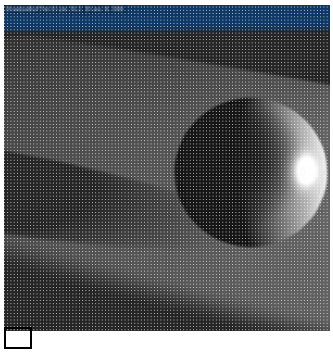
The following settings are common to all buffered shadow generation method.

Bias

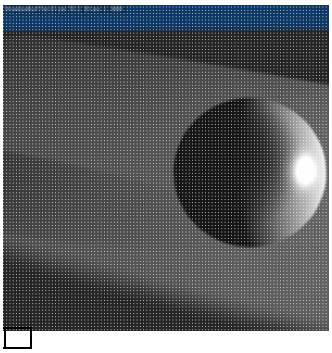
The Bias numeric field can have a value between **0.001** and **5.0**. Bias is used to add a slight offset distance between an object and the shadows cast by it. This is sometimes required because of inaccuracies in the calculation which determines weather an area of an object is in shadow or not. Making the Bias value smaller results in the distance between the object and its shadow being smaller. If the Bias value is too small, an object can get artifacts, which can appear as lines and interference patterns on objects. This problem is usually called â← self shadowingâ← , and can usually be fixed by increasing the Bias value, which exists for that purpose! Other methods for correcting self shadowing include increasing the size of the ShadowBufferSize or using a different buffer shadow calculation method such as Classic-Halfway or Irregular.


Bias: **0.001**.


Bias: **0.1**.



Bias: **0.5**.

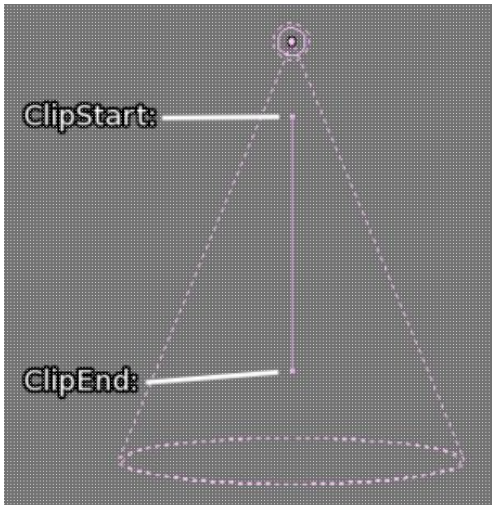


Bias: **1.0**.

Buffered shadow examples for various Bias values (ShadowBufferSize: **512**).

The images above show the affect of different Bias values. With a Bias of **0.001**, the scene is full of self shadowing interferences, but the shadow coming from the back of the sphere is very close to the it. With Bias at **0.1**, most of the self shadowing interferences have been eliminated (apart from small areas on the sphere), but the start of the shadow point has moved slightly to the left of the sphere. With Bias values at **0.5** and **1.0**, the shadow start point moves even further away from the sphere and there is no self shadowing interference.
Self shadowing interference tends to affect curved surfaces more than flat ones, meaning that if your scene has a lot of curved surfaces it may be necessary to increase the Bias value or ShadowBufferSize value.
Having overly large Bias values not only places shadows further away from their casting objects, but can also cause objects that are very small to not cast any shadow at all. At that point altering Bias, ShadowBufferSize or SpotSi values, among other things, may be required to fix the problem.

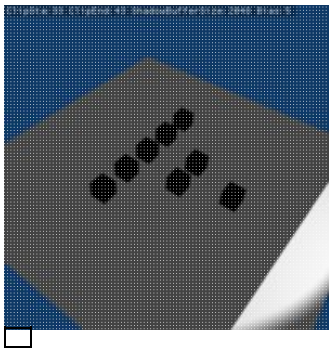
Finer Bias tuning
You can now refine the Bias value independently for each material, using the LBias slider (Shaders panel, Material sub-context, Shading context, F5). This value is a factor by which the Bias value of each Spot buffered shadows lamp is multiplied, each time its light hits an object using this material. The **0.0** and **1.0** values are equivalent â← Â they do not alter the lampâ← s Bias original value.



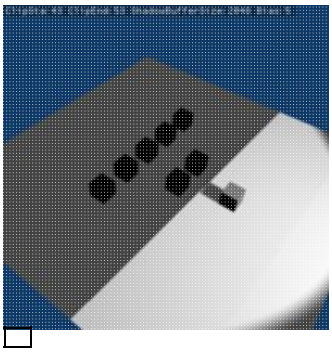
Screenshot showing shadow ClipSta and ClipEnd points line.

ClipSta & ClipEnd

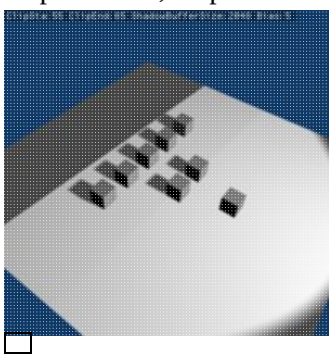
When a Spot light with buffered shadows is added to a scene, an extra line appears on the Spot 3D view representation, see (*Screenshot showing shadow ClipSta and ClipEnd points line*).
The start point of the line represents ClipStaâ← s value and the end of the line represents ClipEndâ← s value. ClipSta can have a value between **0.1** and **1000.0**, and ClipEnd, between **1.0** and **5000.0**. Both values are represented in Blender Units.
ClipStart indicates the point after which buffered shadows can be present within the Spot light area. Any shadow which could be present before this point is ignored and no shadow will be generated.
ClipEnd indicates the point after which buffered shadows will not be generated within the Spot light area. Any shadow which could be present after this point is ignored and no shadow will be generated.



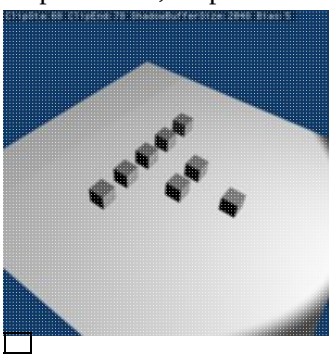
ClipSta: **37.0**, ClipEnd: **47.0**.



ClipSta: **43.0**, ClipEnd: **53.0**.



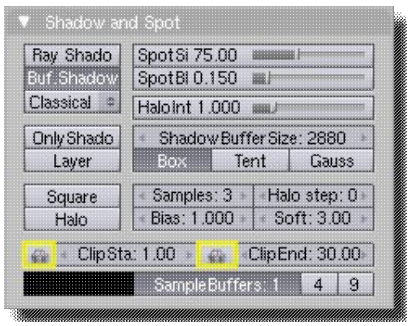
ClipSta: **55.0**, ClipEnd: **65.0**.



ClipSta: **68.0**, ClipEnd: **78.0**.

Buffered shadow examples for various ClipSta and ClipEnd values (ShadowBufferSize: **2048**).

The area between ClipSta and ClipEnd will be capable of having buffered shadows generated. Altering the ClipSta and ClipEnd values helps in controlling where shadows can be generated. Altering the range between ClipSta and ClipEnd can help speed up rendering, save memory and make the resultant shadows more accurate.
When using a Spot lamp with buffered shadows, to maintain or increase quality of generated shadows, it is helpful to adjust the ClipSta and ClipEnd such that their values closely bound around the areas which they want to have shadows generated at. Minimizing the range between ClipSta and ClipEnd, minimizes the area shadows are computed in and therefore helps increase shadow quality in the more restricted area.



Screenshot showing automatic ClipSta and ClipEnd buttons, highlighted in yellow.

Automatic ClipStart & ClipEnd

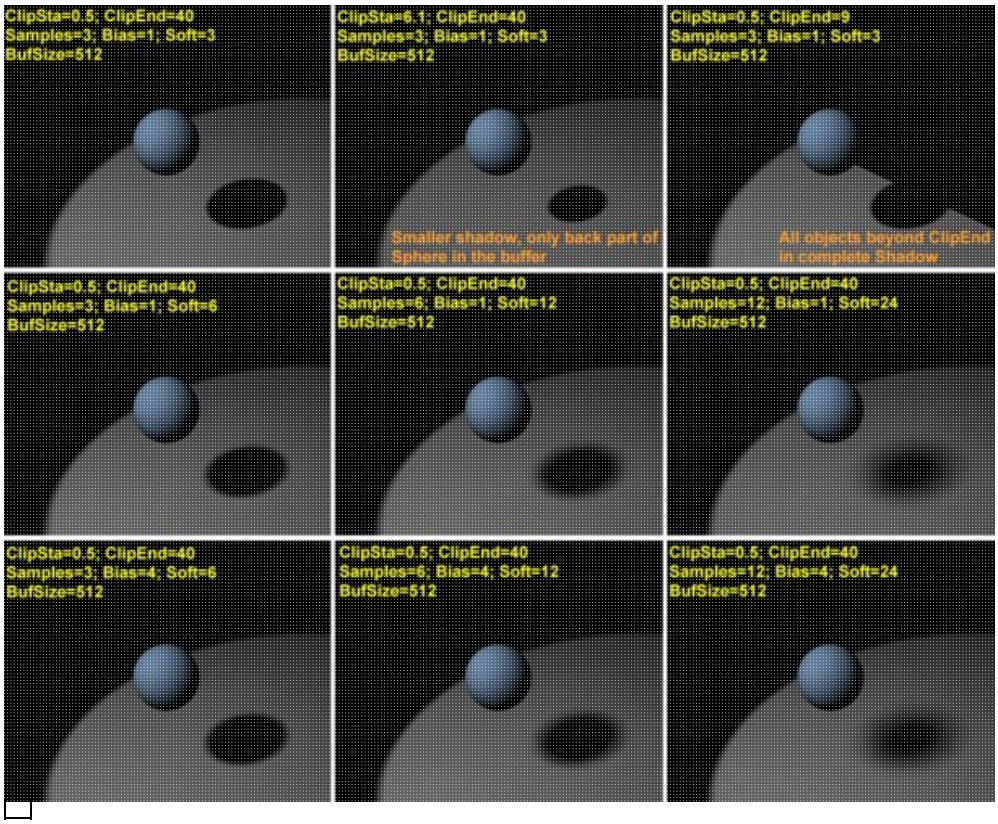
As well as manually setting ClipSta and ClipEnd fields to control when buffered shadows start and end, it is also possible to have Blender pick the best value independently for each ClipSta and ClipEnd field.
Blender does this by looking at where the visible vertices are when viewed from the Spot lamp position.

Note

I am not really clear on how SampleBuffers and Irregular shadows buffers work so if anyone has more info and came make thing clearer either contact me and I will update the page, or update it yourself â← Â the same goes for anything else you think should be changed on this page. Terrywallwork - 3 Oct 2008, Mont29 14:11, 2 May 2009 (UTC).

Examples

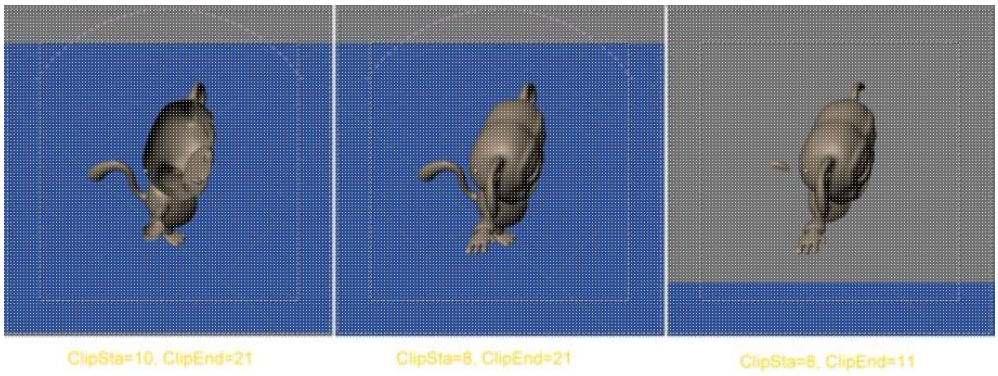
Below are various examples using a solid sphere with shadows cast against a plane object. Note how the different settings affect the shadow that is cast. You get everything from sharp shadows to very soft shadows.



Spot light buffered shadow examples.

Hints

Any object in Blender can act as a camera in the 3D view. Hence you can select the Spot light and switch to a view from its perspective by pressing Ctrl0 NumPad. What you would see, in shaded mode, is shown in (*Spot light clipping tweak*).



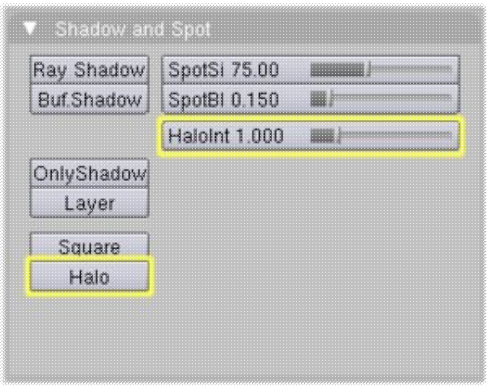
Spot light clipping tweak.

The left frame shows ClipSta too high, the center has it correct, and the right shows ClipEnd too low. All object(s) nearer to the Spot light than ClipSta and further from ClipEnd is(are) not shown at all. Hence you can fine tune these values by verifying that all shadow *casting* objects are visible.

See Also

- Shadows
- Common lamp shadows options
- Spot light source
- Spot raytraced shadows
- Spot halos

Spot Volumetric Effects



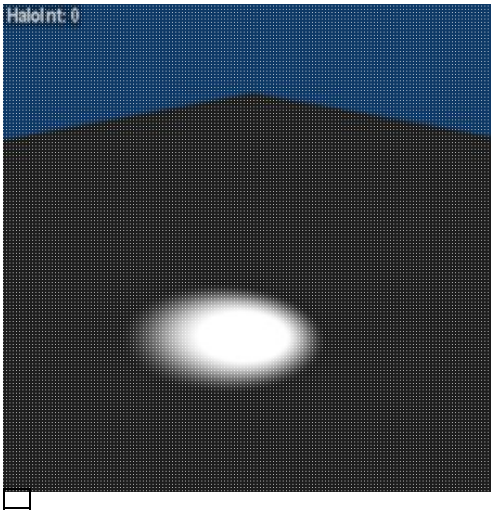
The Shadow and Spot panel for a Spot lamp, with halo controls highlighted.
Spot lights also can produce â← volumetricâ← effects. See this page for more information about what it means.

Halo

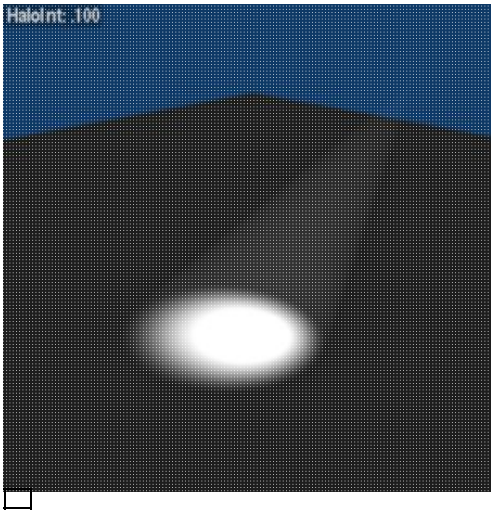
The Halo button allows a Spot lamp to have a volumetric effect applied to it. This button must be active if the volumetric effect is to be visible. Note that if you are using buffered shadows, you have extra options described in the â← Spot Buffered Shadowsâ← page.

HaloInt

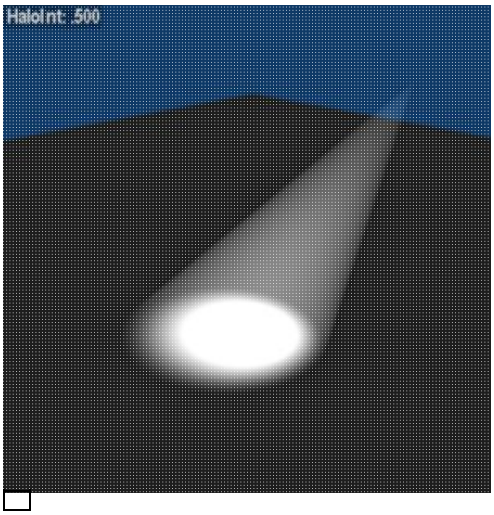
The HaloInt (â← Halo Intensityâ←) slider controls how intense/dense the volumetric effect is, that is generated from the light source. The HaloInt value has a range of between **0.0** and **5.0** â← Â by default, this value is ignored, as the Halo button is not active. To make HaloInt have an effect, make sure the Halo button is active. The lower the value of the HaloInt slider the less visible the volumetric effect is, while higher HaloInt values give a much more noticeable and dense volumetric effect.



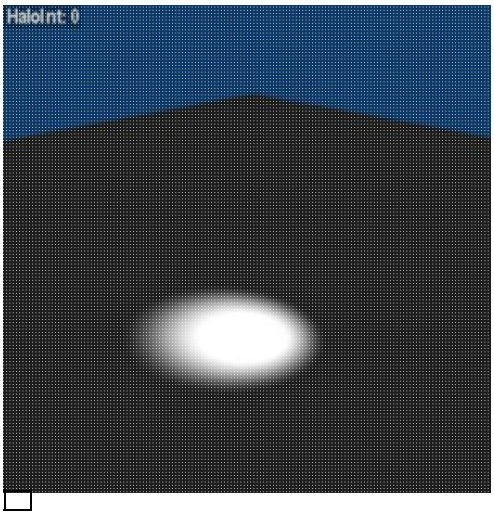
HaloInt: **0.0**.



HaloInt: **1.0**.



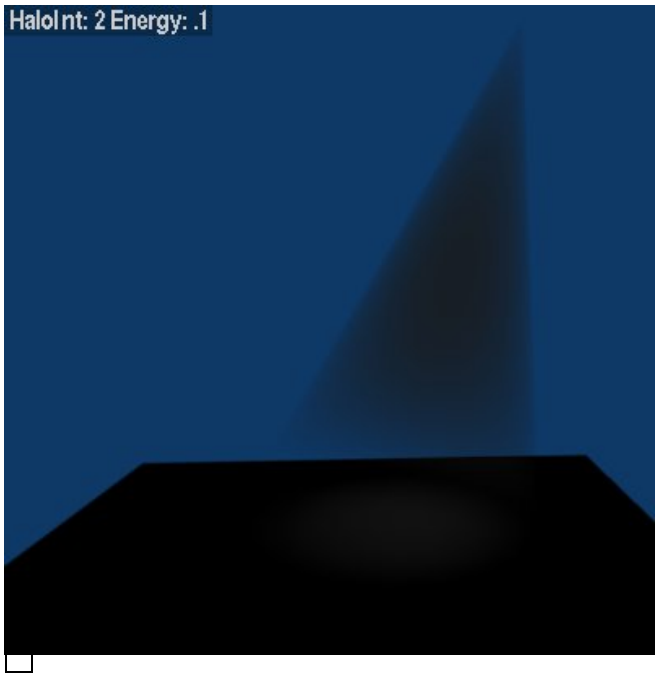
HaloInt: **0.5'**.



Animation showing the effect of different HaloInt values, click to see.

Volumetric halo examples for various HaloInt values.

Blender only simulates volumetric lighting in Spot lamps when using its internal renderer. This can lead to some strange results for certain combinations of settings for lightâ← s Energy and HaloInt. For example having a Spot light with null or very low light Energy settings but a very high HaloInt setting can result in a dark/black halo, which would not happen in the real world. Just be aware of this possibility when using halos with the internal renderer.



Dark halo with low Energy light.

Note
The halo effect can be greatly enhanced when using buffered shadows: when the Halo Step is not null, they can create â← volumetric shadowsâ← . See the page about Spot buffered shadows for more information.

See Also

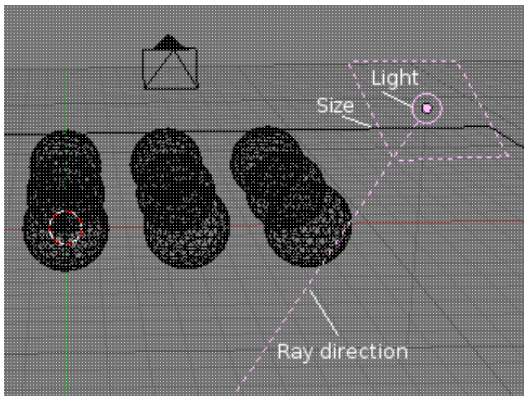
- Shadows
- Spot light source
- Spot buffered shadows

Area

Mode: All modes

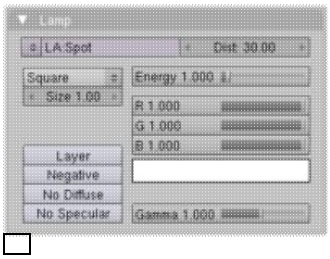
Panel: Lamp (Shading context, Lamp sub-context, F5)

Description

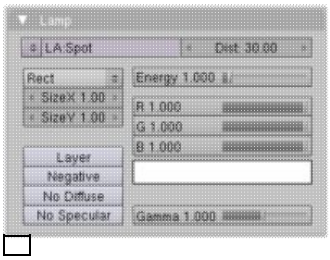


Area lamp.
The Area lamp simulates light originating from a surface (or surface-like) emitter, for example, a TV screen, your supermarketâ← s neons, a window or a cloudy sky are just a few types. The area lamp produces shadows with soft borders by sampling a lamp along a grid the size of which is defined by the user. This is in direct contrast to point-like artificial lights which produce sharp borders.

Options



Square Area lamp.



Rect Area lamp.
Area lightâ← s Lamp panel.

Common options

The Dist, Energy and Color settings are common to most types of lamps, and are described here.
Note that the Dist setting is much more sensitive and important for Area lamps than for others, usually any objects within the range of Dist will be blown out and overexposed. For best results, set the Dist

to just below the distance to the object that you want to illuminate.
The Layer, Negative, No Diffuse and No Specular settings control what the lamp affect, as described in this page.

Gamma
Amount to gamma correct the brightness of illumination. Higher values give more contrast and shorter falloff.

The Area lamp doesnâ← t have light falloff settings. It uses an â← inverse quadraticâ← attenuation law (*TODO: is this true?*). The only way to control its falloff is to use the Dist and/or Gamma settings.

It replaces the light attenuation controls (falloff type and Sphere button) with Shape and Size controls. The first one lets you choose the shape of the area and the second the size of the shape (in Blender Units).

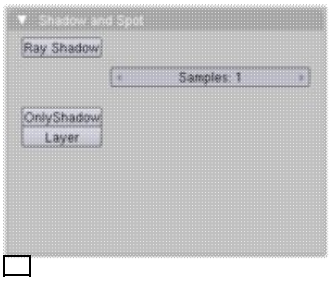
Square
Emit light from a square area.
Size
The length of the squareâ← s edge.

Rect
Emit light from a rectangular area.
SizeX
The rectangleâ← s horizontal width.
SizeY
The rectangleâ← s vertical height.

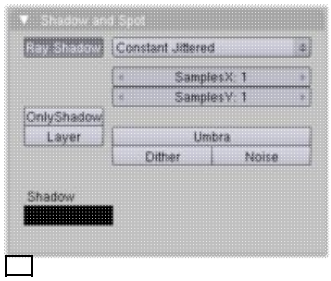
Shape Tips
Choosing the appropriate shape for your Area light will enhance the believability of your scene. For example, you may have an indoor scene and would like to simulate light entering through a window. You could place a Rect area lamp in a window (vertical) or from neons (horizontal) with proper ratios for SizeX and SizeY. For the simulation of the light emitted by a TV-screen a vertical Square area lamp would be better in most cases.

â← Shadow and Spotâ← Panel

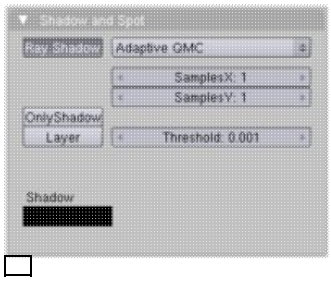
When an Area light source is selected, the Shadow and Spot panel has the following default layout:



Square Area without ray shadows.



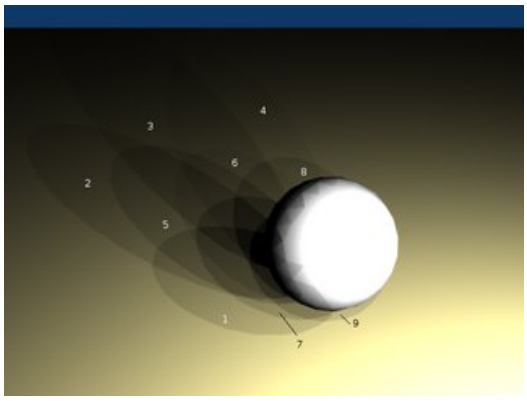
Rect Area with ray shadows, Constant Jittered sample generator.



Rect Area with ray shadows, Adaptive QMC sample generator.

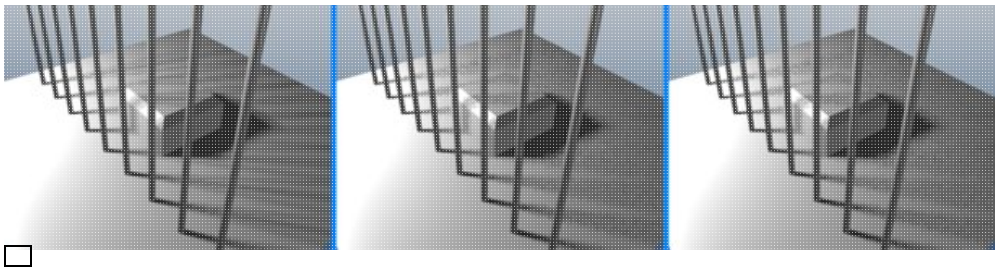
The Shadow and Spot panel when Area light source is selected.
As its only purpose for this lamp is about raytraced shadows, it is described in details here.

Examples



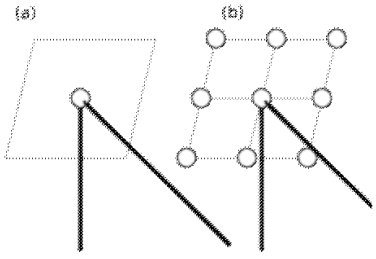
Render example.
In (*Render example*), only one sphere is visible in order to emphasize the shadows created by the Area light, with Constant Jittered sample generator. Here the Samples has been set to **3** which will generate 3Â Ã Â 3Â =Â 9 shadows. In addition, the Size of the square has been made relatively large (**30**) in order to exaggerate the shadows displaced from one another; the numbers are marked in an arbitrary order. Think of the Size as pushing the lights away from each other in the plane of the square.

In the second example below, using the Constant Jittered sample generator, the first image on the left shows a Samples setting of **2** which generates four lights and hence four shadows. You can clearly see the shadows, but if you stand back far enough from the image it will appear as a single â← softâ← shadow. This can be improved by enabling Dithering and improved further by enabling Noise.



Samples to **2**: Dither only, Noise only, and Dither plus Noise.
You can find more useful and complete examples about lamps in the â← lighting rigsâ← page.

Technical Details



Principles behind the Area light.
The (*Principles behind the Area light*) picture helps to understand how the soft shadows are simulated.

(a) is the Area light as defined in Blender. If its shape is Square, then the softness of the shadow is defined by the number of light Samples in each direction of the shape. For example, (b) illustrates the equivalent case of an Area light (Square shape), with Samples set at **3** on the Shadow and Spot panel.

The Area lamp is then considered as a grid with a resolution of three in each direction, and with a light dupliverted at each node for a total of nine lights.

In case (a), the energy (E) is E/1, and in case (b), the energy of each individual pseudo-light is equal to E/(Nbr of lights). Each pseudo-light produces a faint shadow (proportional to its energy), and the overlay of the shadows produces the soft shadow (it is darker where the individual shadows overlap, and lighter everywhere else).

Hints

You will note that changing the Size parameter of your area lamp doesnâ← t affect the lighting intensity of your scene. On the other hand, rescaling the lamp using the S in the 3D View could dramatically increase or decrease the lighting intensity of the scene. This behavior has been coded this way so that you can fine tune all your light settings and then decide to scale up (or down) the whole scene without suffering from a drastic change in the lighting intensity. If you only want to change the dimensions of your Area lamp, without messing with its lighting intensity, you are strongly encouraged to use the Size button(s) instead.

With equal Energy and Dist values, an area lamp and a regular lamp will not light the scene with the same intensity. The area lamp will have a tendency to â← blow outâ← the highlights, but this can be corrected using the Exp slider in the World sub-context.

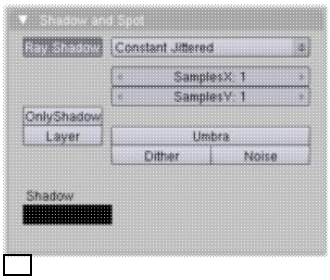
If your computer isnâ← t very fast, when using the Constant Jittered sample generator method, you could find it useful to set a low Samples value (like **2**) and activate Dither and/or Noise in order to simulate slightly softer shadows. However, these results will never be better than the same lighting with high Samples values.

See Also

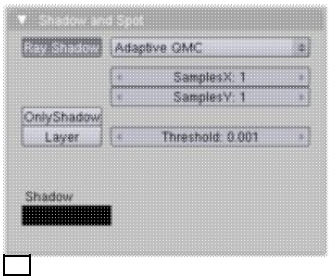
- Shadows
- Area raytraced shadows
- Lighting rigs

Area Raytraced Shadows

The Area light source can only cast raytraced shadows. It shares with other lamp types common shadows options described here.



Rect Area, Constant Jittered sample generator.



Rect Area, Adaptive QMC sample generator.

The Shadow and Spot panel of the Area lamp when ray shadows are enabled.

Raytraced Shadows

The raytraced shadows settings of this lamp are mostly shared with other ones, as described here. However, there are some specificities with this lamp, which are detailed bellow:

Shadow Sample Generator Type

Constant Jittered
: The Area lamp has a third sample generator method, Constant Jittered (I think itâ← s its original one), which is more like simulating an array of lights. It has the same options as the old one: Umbra, Dither and Noise.

Samples
: This have the same role as with other lamps, but when using a Rect Area lamp, you have two samples settings: SamplesX and SamplesY, for the two axes of the area plane!
Note also that when using Constant Jittered sample generator method, this is more or less equivalent to the number of virtual lamps in the area. With QMC sample generator methods, it behave similarly as with Lamp or Spot lamps.

The Soft Size parameter has disappeared.

The following three parameters are only available when using Constant Jittered sample generator method, and are intended to artificially boost the â← softâ← shadow effect, with possible loss in quality:

Umbra
: Emphasizes the intensity of shadows in the area fully within the shadow rays. The light transition between fully shadowed areas and fully lit areas changes more quickly (i.e. a sharp shadow gradient). You need Samples values equal to or greater than **2** to see any influence of this button.

Dither
: Applies a sampling over the borders of the shadows, in a similar same way anti-aliasing is applied by the OSA button on the borders of an object. It artificially softens the borders of shadows; when Samples is set very low, you can expect poor results, so Dither is better used with medium Samples values. It is not useful at all with high Samples values, as the borders will appear already soft.

Noise

Adds noise to break up the edges of solid shadow samples, offsetting them from each other in a pseudo-random way. Once again, this option is not very useful when you use high Samples values where the drawback is that Noise generates quite visible graininess.

See Also

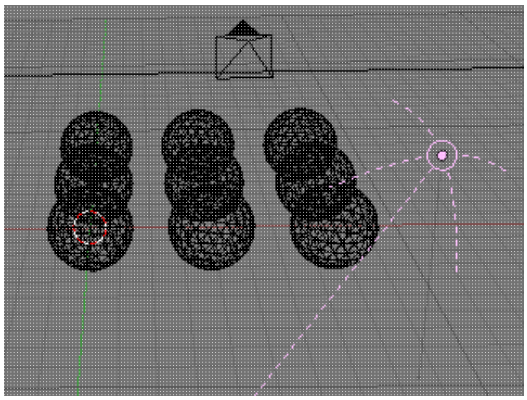
- Shadows
- Common lamp shadows options
- Common lamp ray shadows options
- Area light source

Hemi

Mode: All modes

Panel: Lamp (Shading context, Lamp sub-context, F5)

Description

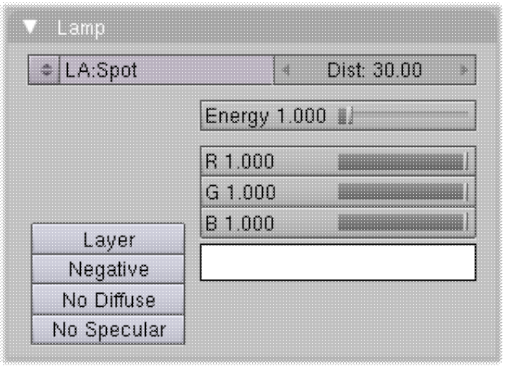


Hemi light.
The Hemi lamp provides light from the direction of a 180Â° hemisphere, designed to simulate the light coming from a heavily clouded or otherwise uniform sky. In other words, it is a light which is shed, uniformly, by a glowing dome surrounding the scene (*Hemi light conceptual scheme*).

Similar to the Sun lamp, the Hemiâ← s location is unimportant, while its orientation is key.

The Hemi lamp is represented with four arcs, visualizing the orientation of the hemispherical dome, and a dashed line representing the direction in which the maximum energy is radiated, the inside of the hemisphere.

Options



Lamp panel.

Common options

The Dist, Energy and Color settings are common to most types of lamps, and are described here. However, note that as this lamp always uses a constant light falloff, the Dist setting doesnâ← t have any effect in practice!

Todo
Is it true that for Sun and Hemi lamps, Dist has no effect at all? -- Mont29 14:13, 2 May 2009 (UTC).

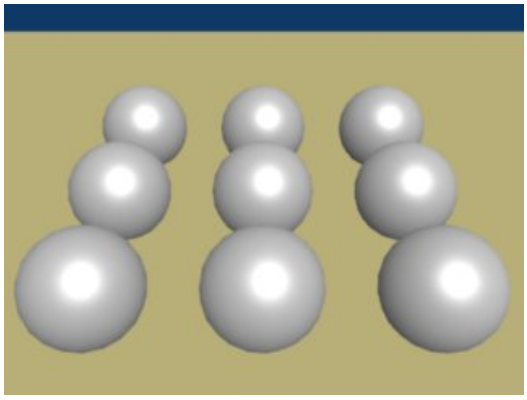
> The Layer, Negative, No Diffuse and No Specular settings control what the lamp affect, as described in this page.

The Hemi lamp has no light falloff settings: it always uses a constant attenuation (i.e. no attenuation!).

â← Shadow and Spotâ← Panel

When a Hemi light source is selected, the Shadow and Spot panel is completely empty: this lamp is the only one in Blender which cannot cast any shadow.

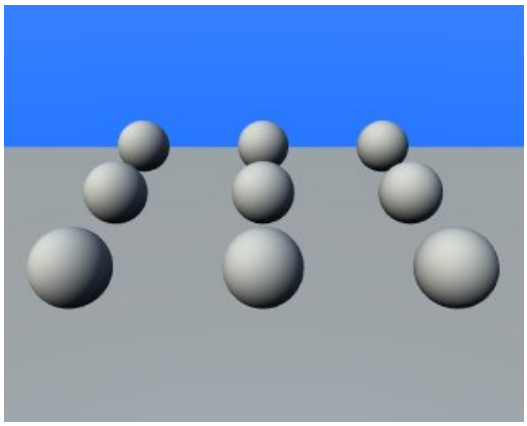
Examples



Render example.
The results of a Hemi light for the nine spheres set are shown in (*Render example*). Notice how the spheres are lit more completely around and to the backside, and the beige ground appears to be bright. The softness of the Hemi light in comparison to the Sun light is evident.

You can find more useful and complete examples about lamps in the â← lighting rigsâ← page.

Outdoor Light



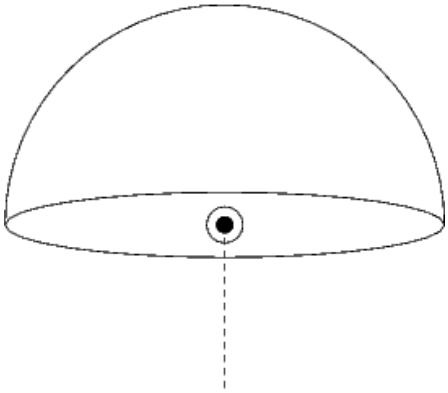
Outdoor light example.

To achieve outdoor lighting you can use both a Sun light, say an Energy of **1.0**, with a warm yellow/orange tint and shadowing enabled, plus a weaker bluish Hemi light faking the light coming from every point of a clear blue sky.

(*Outdoor light example*) shows an example with relative parameters. The configuration is as:

- Sun light: Energy=**1.0**; RGB=(**1.0**,**0.95**,**0.8**); Sun direction in a polar reference is (**135Â°**,**135Â°**).
- Hemi light: Energy=**0.5**; RGB=(**0.64**,**0.78**,**1.0**); pointing down.

Technical Details



Hemi light conceptual scheme.

See Also

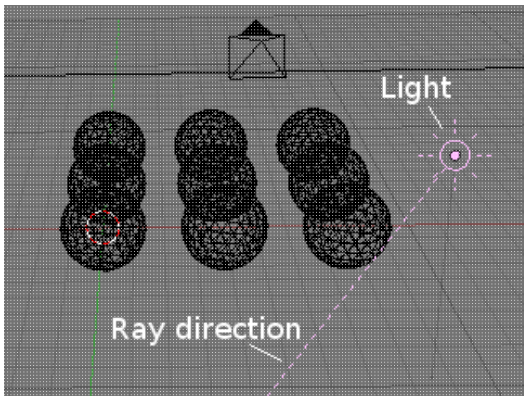
- Lighting rigs

Sun

Mode: All modes

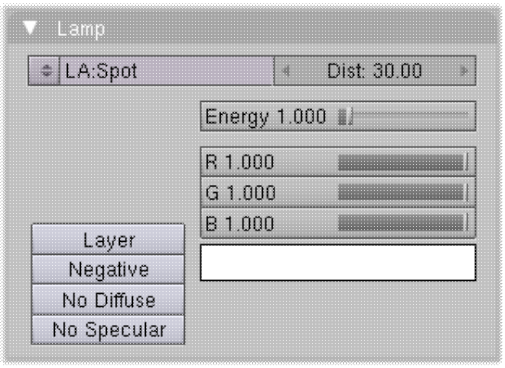
Panel: Lamp (Shading context, Lamp sub-context, F5)

Description



Sun light.
A Sun lamp provides light of constant intensity emitted in a single direction. In the 3D view, the Sun light is represented by an encircled black dot with rays emitting from it, plus a dashed line indicating the direction of the light. This direction can be changed by rotating the Sun lamp, as any other object, but because the light is emitted in a constant direction, the location of a Sun lamp does not affect the rendered result (unless you use the â← sky & atmosphereâ← optionâ←¦).

Options



Lamp panel.

Common options

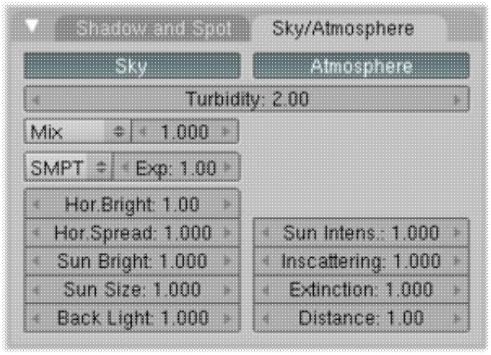
The Dist, Energy and Color settings are common to most types of lamps, and are described here. However, note that as this lamp always uses a constant light falloff, the Dist setting doesnâ← t have any effect in practice!

Todo
Is it true that for Sun and Hemi lamps, Dist has no effect at all? -- Mont29 14:19, 20 June 2009 (UTC).

The Layer, Negative, No Diffuse and No Specular settings control what the lamp affect, as described in this page.

The Sun lamp has no light falloff settings: it always uses a constant attenuation (i.e. no attenuation!).

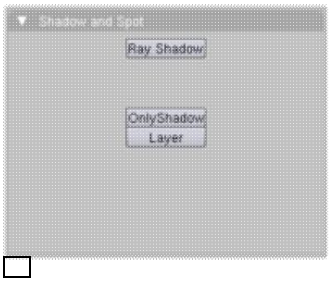
â← Sky/Atmosphereâ← Panel



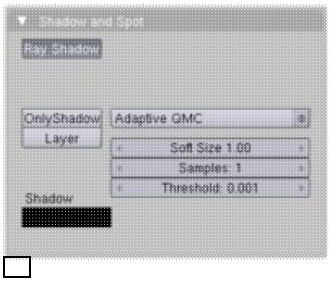
The Sky/Atmosphere panel of the Sun light source.
Various settings for the appearance of the sun in the sky, and the atmosphere through which is shines, are available. For details, see this page.

â← Shadow and Spotâ← Panel

When a Sun light source is selected the Shadow and Spot panel has the following default layout:



Without ray shadows.

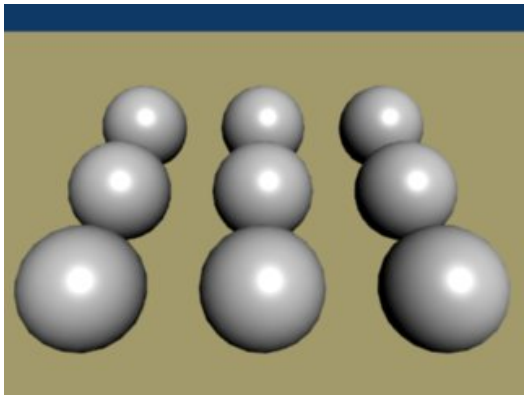


With ray shadows, Adaptive QMC sample generator.

The Shadow and Spot panel when Sun light source is selected.

As its only purpose for this lamp is about raytraced shadows, it is described in details here.

Examples



Render example.
In this rendering the light comes from a constant direction and has a uniform intensity. Notice also, that the white specular highlight on each sphere is in the exact same place (discounting of course the perspective effect from the cameraâ← s position being so close to the spheres). This means the actual *location* of the Sun light itself is not important (for its lighting).

You can find more useful and complete examples about lamps in the â← lighting rigsâ← page.

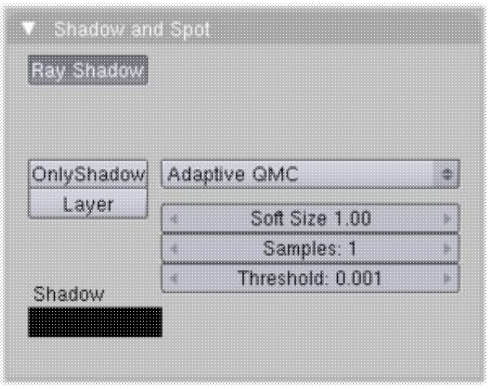
Hints

A Sun lamp can be very handy for a uniform clear day-light open-space illumination.

See Also

- Sky & Atmosphere
- Shadows
- Sun raytraced shadows
- Lighting rigs

Sun Raytraced Shadows



The Shadow and Spot panel for a Sun lamp, with raytraced shadows activated.
The Sun light source can only cast raytraced shadows. It shares with other lamp types the same common shadowing options, described here.

The raytraced shadows settings of this lamp are shared with other ones, and are described here.

See Also

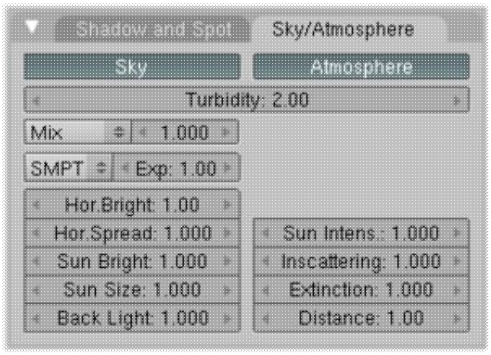
- Shadows
- Common lamp shadows options
- Common lamp ray shadows options
- Sun light source

Sun: Sky & Atmosphere

Mode: All modes (Sun lamp)

Panel: Sky/Atmosphere (Shading context, Lamp sub-context, F5)

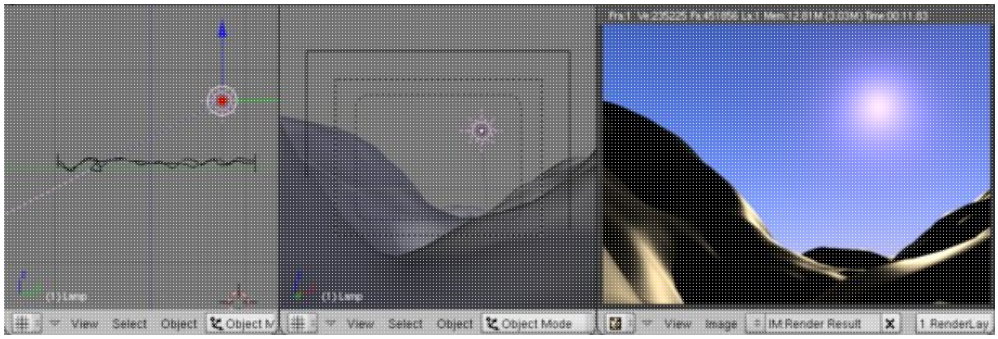
Description



The Sky/Atmosphere panel of the Sun light source.
This panel allows you to enable a new effect that simulates various properties of real sky and atmosphere: the scattering of the sun light as it crosses the kilometers of air overhead. For example, when the Sun is high, the sky is blue (and the horizon, somewhat whitish). When the Sun is near the horizon, the sky is dark blue/purple, and the horizon turns orange. The dispersion of the atmosphere is also more visible, when it is a bit foggy: the more away is an object, the more â← fadedâ← in light grey it isâ←¦ Go out in the countryside, a nice hot day, you will seeâ←¦

To enable this effect, you have to use a Sun light source. If, as usual, the *position* of the lamp has no importance, its *rotation* is crucial: it determines which hour it is! As a start point, you should reset rotation of your Sun (with AltR, or typing **0** in each of the three fields RotX/RotY/RotZ in the Transform Properties panel â← Â N). This way, youâ← ll have a nice mid-day sun (in the tropics!).

Now, there are two important angles for the Sky/Atmosphere effect: the â← inclinationâ← angle (between light direction and X-Y plane), which determine the â← hourâ← of the day (as you might expect, default rotation â← Â straight downÂ â← is â← mid-dayâ← , a light pointing straight up is â← mid-nightâ← , and so onâ←¦). And the rotation around the Z axis determines the position of the sun around the camera.



The dashed â← light lineâ← of the Sun lamp crossing the camera focal point.
In fact, to have a good idea of where the sun is in your world, relative to the camera, you should always try to have in your 3D view the dashed â← light lineâ← of the lamp crossing the center of the camera (its â← focalâ← point), as shown in (*The dashed â← light lineâ← of the Sun lamp crossing the camera focal*

point). This way, in camera point of view (0 NumPad, center window in the example picture), you will see where will be the â← virtualâ← sun created by this effect.

It is important to understand that the *position* of the sun has no importance for the effect: only its *orientation* is relevant. The position might just help you in your scene design.

Options

Turbidity

This is a general parameter that affects sun view, sky and atmosphere, itâ← s an atmosphere parameters that low values describe clear sky, and high values shows more foggy sky. In general, low values give clear, deep blue sky, with â← littleâ← sun; high values give more reddish sky, with a big halo around the sun. Note that this parameter is one which can really modify the â← intensityâ← of the sun lighting. See examples below.

Sky

Sky

This button enables the sky settings: it will create a â← skyâ← , with a â← sunâ← if visible, and mix it with the background as defined in World settings.

Here are its specific controls:

Mix menu

The first drop-down list shows you a menu of various mix methods. The one selected will be used to blend the sky and sun with the background defined in the World settings. The mixing methods are the same as described e.g. in the Mix compositing node page. And the numeric field at the right of the list allows you to control how much the sky and sun effect is applied to the World background.

Color space menu

This drop-down list allows you to select which coloor space the effect uses, with following choices:

- ◊ CIE
- ◊ REC709
- ◊ SMPTE

Note: I do not understand very well this controlâ←¦ But it surely affects the colors of the effect!

Exp

This numeric field allows you to modify the exposition of the rendered Sky and Sun (**0.0** for no correction).

Hor.Bright

For â← Horizontal Brightnessâ← , controls brightness of horizon colors. Its value should be in range **0.0** to **10.0**, values near zero means no horizontal brightness, and large values for this parameter increase horizon brightness. See examples below.

Hor.Spread

For â← Horizontal Spreadâ← , controls spread of light at horizon. Its value should be in range **0.0** to **10.0**, low values in the range result in less spread of light at horizon, and high values in the range result all horizon light spread in sky.

Sun Bright
: For â← Sun Brightnessâ← , controls sun brightness! Its value should be in range **0.0** to **10.0**, with low values the sky has no sun and with high values the sky only has sun.

Sun Size
: Controls size of sun. Its values should be in range **0.0** to **10.0**, but note that low values result in large sun size, and high values result in small sun size. Note that the overall brightness of the sun remains constant (set by Sun Bright), so the larger the sun (the smaller Sun Size!), the more it â← vanishesâ← in the sky, and *vice-versa*.

Back Light
: For â← Back Scatter Lightâ← , result on sunâ← s color, high values result on more light around the sun. Its values range is **-1.0** to **1.0**. Negative values result on no more light around sun.

Atmosphere

Atmosphere
: This button enables the atmosphere settings. It will not modify the background, but it tries to simulate the effects of an atmosphere: scattering of the sun light in the atmosphere, its attenuation, â←¦

Sun Intensity
: Sets sun intensity. Its values are in range **0.0** to **10.0**. High values result more blue light in far objects.

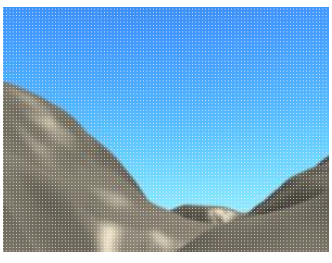
Inscattering
: This factor can be used to decrease the effect of light inscattered into atmosphere between camera and objects in scene. This value should be **1.0** but its change may results in some nice, not realistic, images.

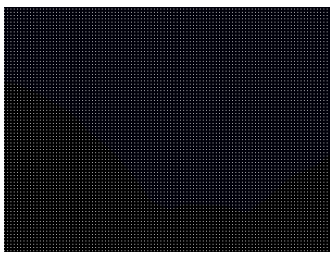
Extinction
: This factor can be use to decrease the effect of extinction light from objects in the scene. Like Inscattering factor, this parameter should be **1.0** but you can change it, low values result in less light extinction. Its value is in range **0.0** to **1.0**.

Distance
: This factor is used to convert Blender units into an understandable unit for atmosphere effect, it starts from **0** and high values result in more yellow light into the scene.

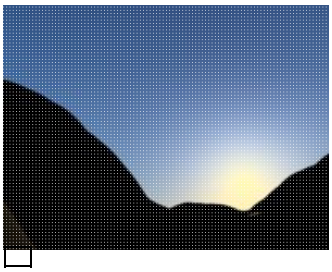
Examples

First, letâ← s see what happens when we modify the orientation of the sun:



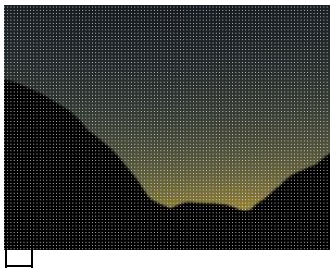


With sun right overhead (mid-day).



Sun slightly above the horizon (start of twilight).

With sun deep â← under the Earthâ← (mid-night).

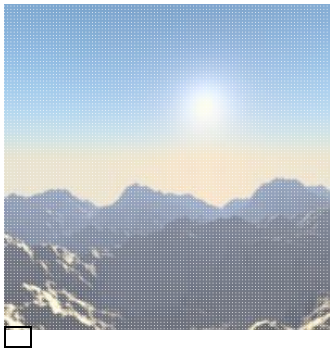


Sun slightly below the horizon (end of twilight).

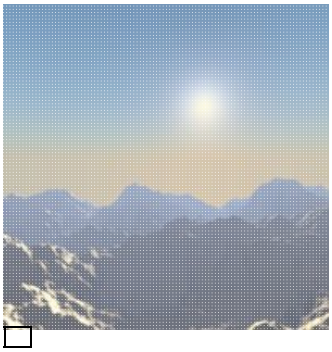
Variations in Sun orientation, Sun Size to **5.0**, all other settings to default.

The .blend file of these examples.

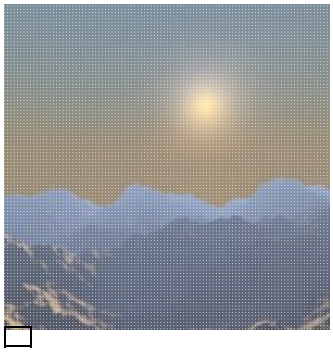
And now, the effects of various settings (examples created with this .blend file):



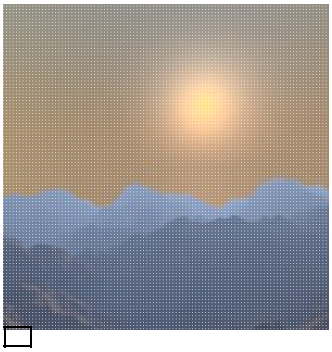
Turbidity: **2.0**.



Turbidity: **2.3**.



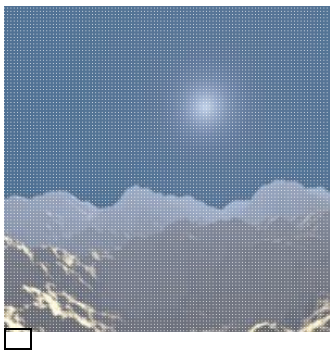
Turbidity: **5.0**.

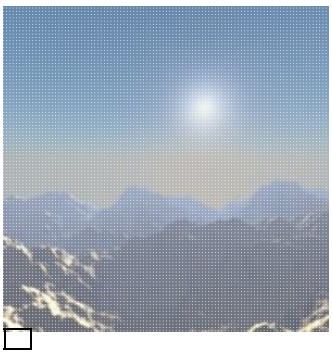


Turbidity: **10.0**.

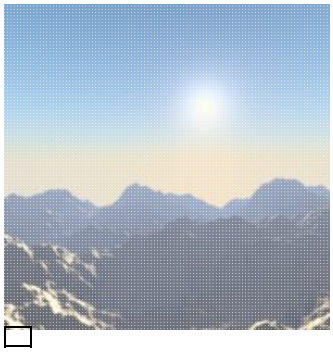
Variations in Turbidity parameter, all other settings to default.

Sky

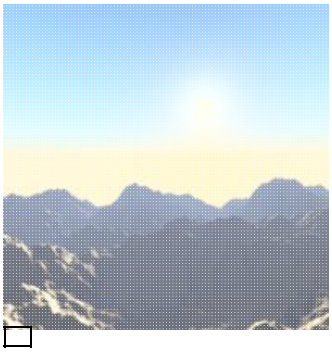




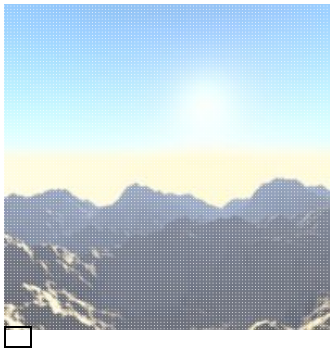
Hor.Bright: **0.0**.



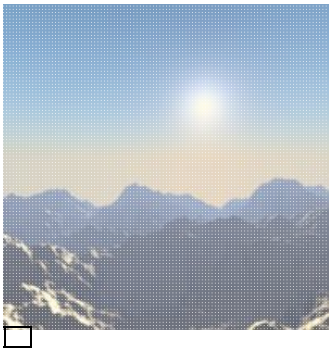
Hor.Bright: **0.85**.



Hor.Bright: **1.04**.



Hor.Bright: **1.13**.

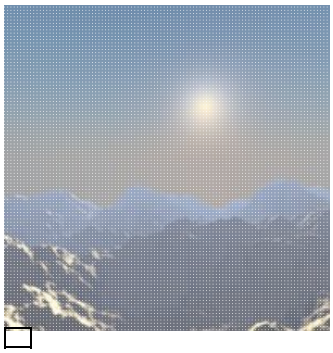


Variations in Hor.Bright parameter, all other settings to default.

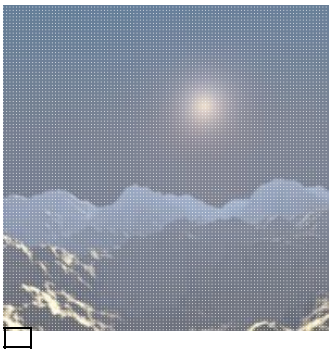
Hor.Spread: **0.7**.

Hor.Spread: **1.2**.

Hor.Spread: **2.2**.

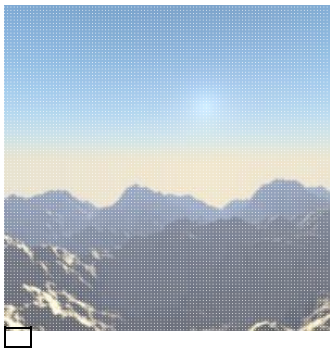


Hor.Spread: **5.0**.

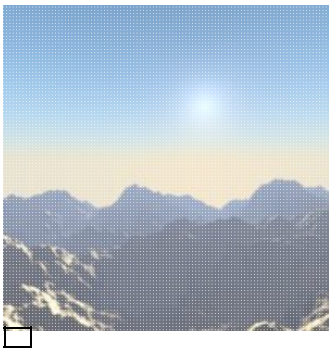


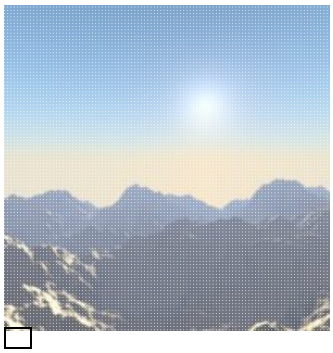
Variations in Hor.Spread parameter, all other settings to default.

Sun Bright: **0.2**.

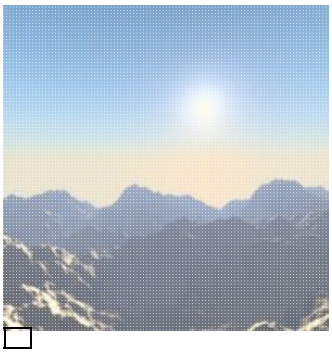


Sun Bright: **0.5**.



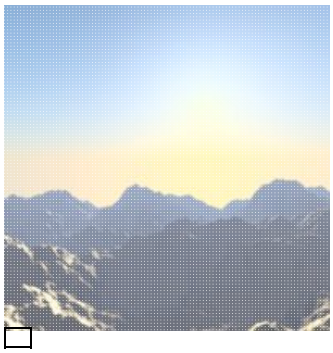


Sun Bright: **0.75**.

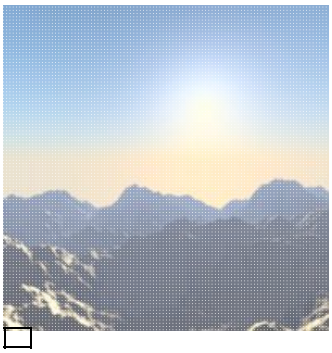


Sun Bright: **1.0**.

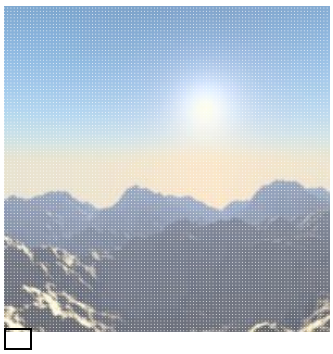
Variations in Sun Bright parameter, all other settings to default.



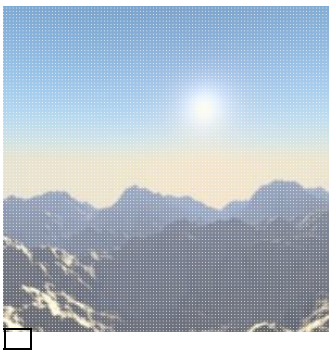
Sun Size: **2.0**.



Sun Size: **4.0**.

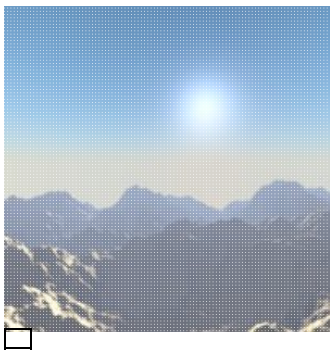


Sun Size: **7.0**.

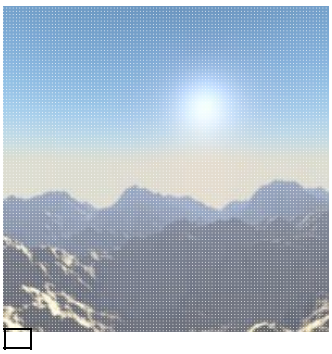


Sun Size: **10.0**.

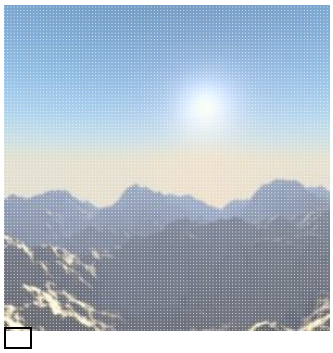
Variations in Sun Size parameter, all other settings to default.



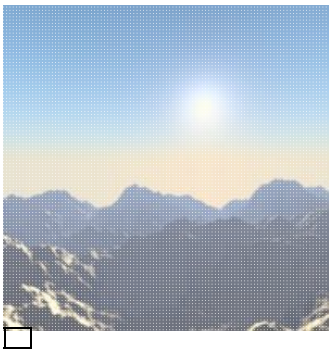
Back Light: **-1.0**.



Back Light: **-0.33**.



Back Light: **0.33**.

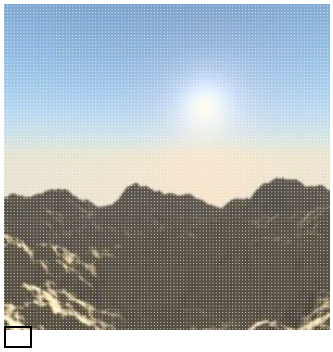


Back Light: **1.0**.

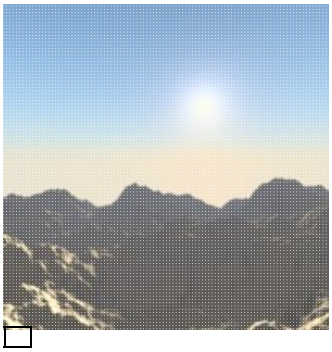
Variations in Back Light parameter, Sun Bright to **2.5**, all other settings to default.

Atmosphere

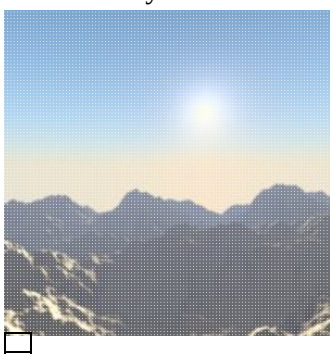
For all renders below, Hor.Bright is set to **0.2**, and Sun Bright to **2.0**.



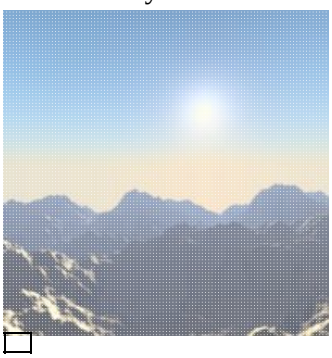
Sun Intensity: **1.0**.



Sun Intensity: **3.33**.

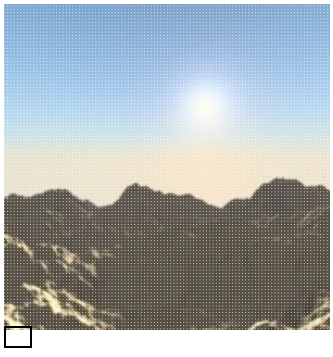


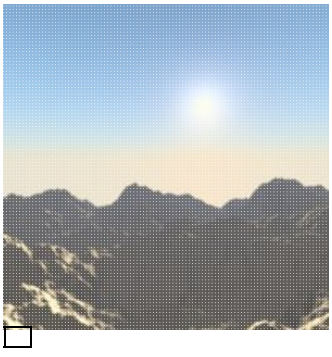
Sun Intensity: **6.66**.



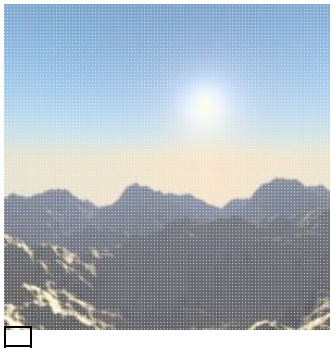
Sun Intensity: **10.0**.

Variations in Sun Intensity parameter, all other settings to default.

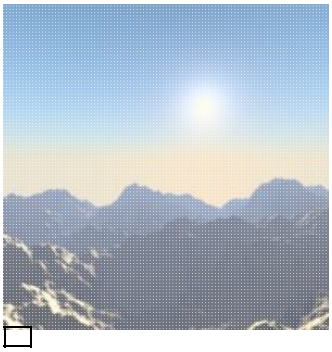




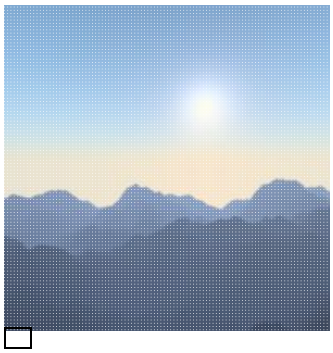
Inscattering: **0.1**.



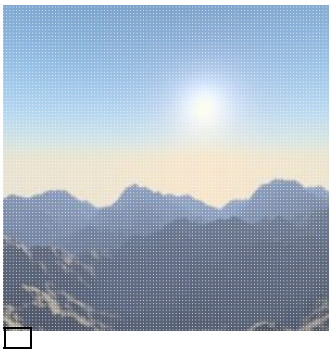
Inscattering: **0.33**.



Inscattering: **0.66**.

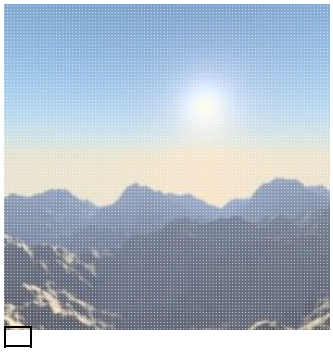


Inscattering: **1.0**.

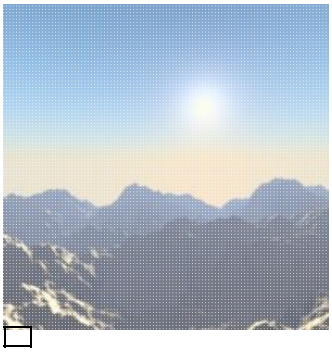


Variations in Inscattering parameter, all other settings to default.

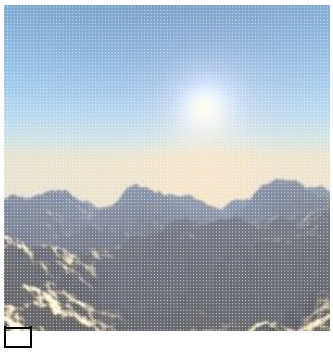
Extinction: **0.0**.



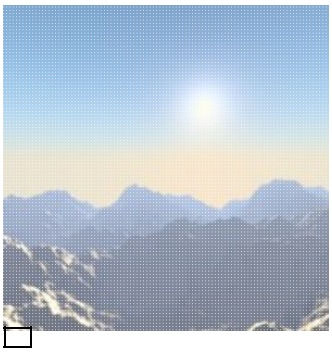
Extinction: **0.33**.



Extinction: **0.66**.



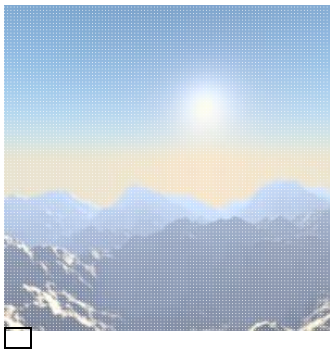
Extinction: **1.0**.



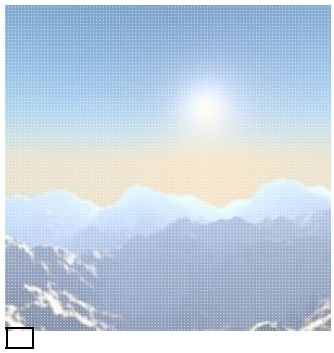
Variations in Extinction parameter, all other settings to default.

Distance: **1.0**.

Distance: **2.0**.



Distance: **3.0**.



Distance: **4.0**.

Variations in Distance parameter, all other settings to default.

Hints and limitations

To always have the Sun pointing at the camera center, you can use a TrackTo constraint on the sun object, with the camera as target, and -Z as the â← Toâ← axis (use either X or Y as â← Upâ← axis). This way, to modify high/position of the sun in the rendered picture, you just have to move it, orientation is automatically handled by the constraint! Of course, if your camera itself is moving, you should also add e.g. a Copy Location constraint to your Sun lamp, with the camera as target â← Â and the Offset option activatedâ←¦ This way, the sun light wonâ← t change as the camera moves around!

If you use the default Add mixing type, you should use a very dark-blue world color, to get correct â← nightsâ← â←¦

This effect works quite well with a Hemi lamp, or some ambient occlusion, to fill-in the Sun shadows.

Atmosphere shading currently works incorrectly in reflections and refractions and is only supported for solid shaded surfaces. This will be addressed in a later release.

See Also

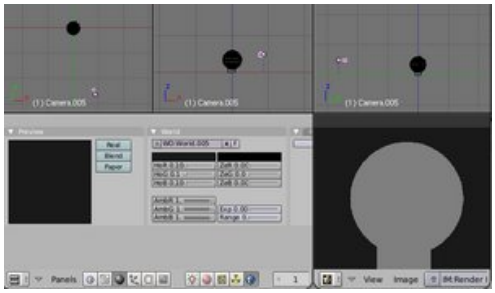
- Sun light
- Lighting rigs
- The release notes about this new feature

Lighting Rigs

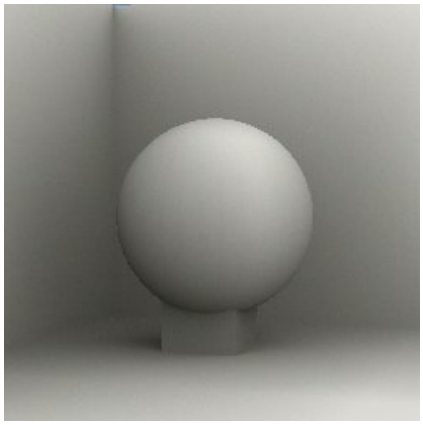
A rig is a standard setup and combination of objects; there can be lighting rigs, or armature rigs, etc. A rig provides a basic setup and allows you to start from a known point and go from there. Different rigs are used for different purposes and emulate different conditions; the rig you start with depends on what you want to convey in your scene. Lighting can be very confusing, and the defaults do not give good results. Further, very small changes can have a dramatic effect on the mood and colors. At major studios, lighting is an entire step and specialty. Well, letâ← s get out of the darkness of confusion and let me en*light*en you.

In all the lighting rigs, the default camera is always positioned **15** degrees off dead-on, about **25 BU** (Blender Units) back and **9 BU** to the side of the subject, at eye level, and uses a long lens (**80mm**). Up close, a **35mm** lens will distort the image. A long lens takes in more of the scene. A dead-on camera angle is too dramatic and frames too wide a scene to take in. So now you know; next time you go to a play, sit off-center and you wonâ← t miss the action happening on the sidelines and will have a greater appreciation for the depth of the set. Anyway, enough about camera angles; this is about lighting.

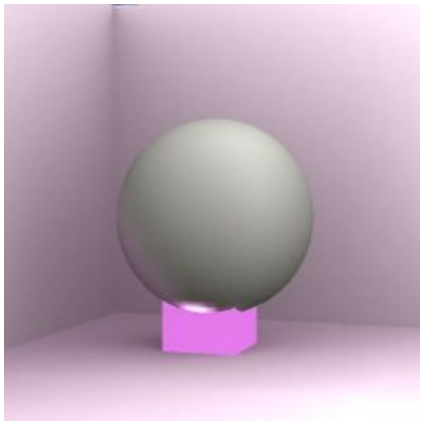
Ambient Only



Ambient lighting only.
In the Shading context (F5), there is a little globe button where you select the World sub-context settings. On the World panel are three sliders AmbR, AmbG, and AmbB, for the color and saturation of ambient light in the world. Ambient light is the scattered light that comes from sunlight being reflected off every surface it hits, hitting your object, and traveling to camera.

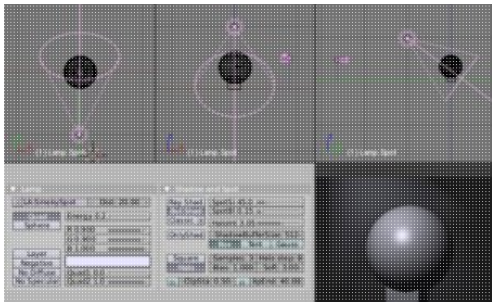


Ambient occlusion.
Ambient light illuminates, in a perfectly balanced, shadeless way, without casting shadows. You can vary the intensity of the ambient light across your scene via ambient occlusion. In the sample .blend, the â← `0 pt AO`â← scene has a light box, Shadow and Ray enabled, and Ambient Occlusion enabled in the World sub-context. The ambient color is a sunny white.



Ambient occlusion with radiosity.
You can add to the ambient light via radiosity. With radiosity, the color of light that is radiated from colored objects is mixed with the ambient light. In the sample .blend, the â← `radiosity`â← scene has Radiosity enabled, and the cube is emitting just a little purple to show the effects. You can mix AO and radiosity as well, as is shown in the example picture to the right and the .blend file in the â← `0 pt AO+Radio`â← scene.

Single Rig



Standard Spot light rig.
The sole, or key, spot light rig provides a dramatic, showy, yet effective illumination of one object or a few objects close together. It is a single Spot light, usually with a hard edge. Halos are enabled in this render to remind you of a smoky nightclub scene. It is placed above and directly in front of the subject; in this case **10 BU** in front and **10 BU** high, just like a stage, it shines down at about a **40** degrees angle. We use quadratic attenuation, energy of **0.2**, a falloff Distance of **20.0** (the light is about **14 BU** from the subject), and a slight coloration of relaxing blue (R:**0.9**, G:**0.9**, B:**1.0**) with a cloud texture mapped to white color at **0.5** Mix. This mixing gives some softness to the light. Itâ← s a narrow spot at **45** degrees, and the halo you can adjust to your liking.

You can make the spot wider by increasing SpotSi and softening the edge by increasing SpotBl, and parent it to the main actor, so that the spot follows him as he moves around. Objects close to the main actor will naturally be more lit and your viewer will pay attention to them.

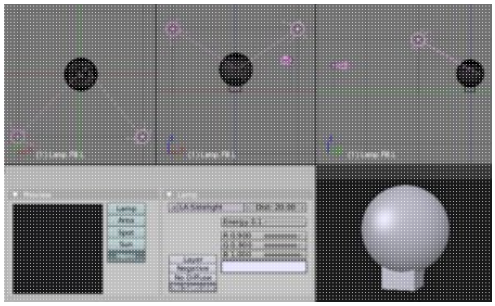
Moving this spot directly overhead and pointing down gives the interrogation effect. At the opposite end of the show-off emotional spectrum is one soft candlelight (Lamp lamp, short falloff Dist, yellow light) placed really up close to the subject, dramatizing the fearful â← lost in the darknessâ← effect.

Somewhere in the macabre spectrum is a hard spot on the floor shining upward. For fun, grab a flashlight, head into the bathroom and close the door. Turn out the light and hold the flashlight under your chin, pointing up. Look in the mirror and turn it on. Ghoulies! Donâ← t blame me for nightmares, and I hope you get the

point: lighting, **even with a single light, varying the intensity, location and direction, changes everything** in a scene.

Use this rig, with Ambient light (and props receiving and being lit by ambient light in their material settings) for scenes that feature one main actor or a product being spotlighted. Do not use this rig for big open spaces or to show all aspects of a model.

Two-Point Rig



Standard two-point light rig.
The two-point lighting rig provides a balanced illumination of an object. Shown to the right are the views of the standard two-point lighting rig. It is called the two-point because there are two points of light. The standard two-point lighting rig provides a balanced illumination of untextured objects hanging out there in 3D space. This rig is used in real studios for lighting a product, especially a glossy one.

Both lights are almost the same but do different things. Both emulate very wide, soft light by being Hemi with a long falloff Distance of **20.0**, which provides even lighting front to back (??? TODO: I thought hemi lamps didnâ← t have any falloff! -- Mont29 14:17, 2 May 2009 (UTC)). Both are tinged blue (R:**0.9**, G:**0.9**, B:**1.0**), and have a white cloud texture mixed at **50%**. If you use a spot light, you will get a shadow. In real life, these lights bounce light off the inside of a silver umbrella.

The setting for the stage left (on your right) light are shown (Shading context, Lamp sub-context, Lamp panel). Notice how we use low Energy to bring out the dimensionality of the sphere; I canâ← t stress that enough. Hard, bright lights actually flatten it and make you squint. Soft lights allow your eye to focus. The left lamp is energy **0.17**, because it gives a little more face to the camera (itâ← s on the same side as the camera, so it more directly lights up the face for the camera), and so we disable specular (No Specular button) so we donâ← t get that shiny forehead or nose.

The lamp on the left however, lets it be known that it is there by enabling specular; specular flare is that bright spot that is off center above midline on the sphere. It functions also as fill, so we turn down its Energy to **0.1**.

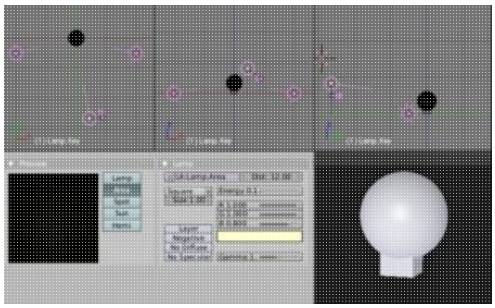
Use this rig to give even illumination of a scene, where there is no main focus. The Hemiâ← s will light up background objects and props, so Ambient is not that important. At the opposite end of the lighting spectrum, two narrow Spot lights at higher power with a hard edge gives a â← This is the Police, come out with your hands upâ← kind of look, as if the subject is caught in the crossfire.

Three-Point Rigs

The standard three-point lighting rig is the most common illumination of objects and scenes bar none. If you want to show off your model, use this rig. As you can see, the untextured unmaterialized sphere seems to come out at you. There are multiple thesis on this rig, and you will use one of two:

- Studio â← used in a real studio to film in front of a green screen or backdrop. Use this rig when you are rendering your CG objects to alpha into the scene so that the lighting on the actors *and* your CG objects is the same.
- Standard â← used in real life to light actors on a set, and gives some backlighting to highlight the sides of actors, making them stand out more and giving them depth.

Studio rig



Studio three-point light rig.
Shown to the right are the â← Studioâ← top, front, and side views of the standard three-point lighting rig. It changes the dynamics of the scene, by making a brighter â← keyâ← light give some highlights to the object, while two side â← fillâ← lights soften the shadows created by the key light.

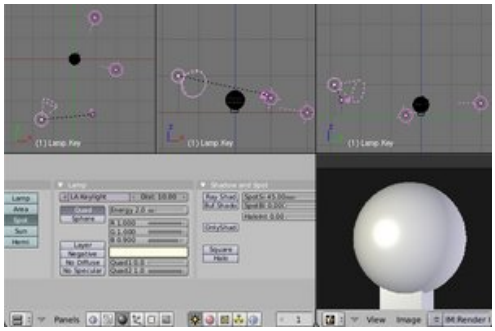
In the studio, use this rig to film a talking head (actor) in front of a green screen, or with multiple people, keeping the key light on the main actor. This rig is also used to light products from all angles, and the side fill lights light up the props.

The key light is the Area light placed slightly above and to the left of the camera. It has an energy of **0.1**, a falloff Dist of **12.0**, is slightly yellow (R:**1.0**, G:**1.0**, B:**0.8**), and it allows the specular to come out. It is about **30 BU** back from the subject, and travels with the camera. A little specular shine lets you know thereâ← s a light there, and that youâ← re not looking at a ghost. In real life, it is a spot with baffles, or blinders, that limit the area of the light.

The two sidelights are reduced to only fill; each of them are Hemi lights placed **20 BU** to the side and **5 BU** in front of the subject, at ground level. Each have an energy of **0.2**, falloff Distance **10.0** and are slightly blue (R:**0.9**, G:**0.9**, B:**1.0**), although I have seen very blue (R:**0.67**, G:**0.71**, B:**0.9**) used effectively. They donâ← t cause a spotshine on the surface by disabling specular (No Specular button), and at ground level, light under the chin or any horizontal surfaces, countering the shadows caused by the key light. Further, a cloud texture is mapped to white at a **50%** mix. This mixing simulates what happens in real life and softens the lighting even further.

Use this rig to give balanced soft lighting that also highlights your main actor or object. It combines the best of both the single rig and the two-point rig, providing balanced illumination and frontal highlights. For a wide scene, you may have to pull the sidelights back to be more positioned like the two-point rig.

Standard Rig

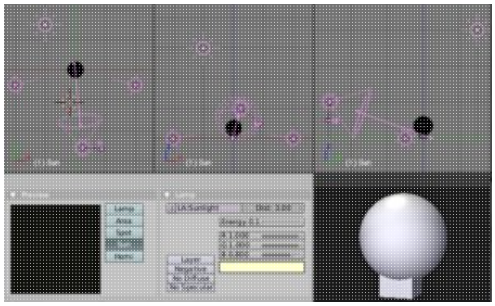


Standard three-point light rig.
Without a curtain in back of your main subject, you have depth to work with. The left fill light has been moved behind the subject (so it is now called a backlight) and is just off-camera, while the right side fill light remains the same. The keylight gives you specular reflection so you can play with specularity and hardness in your objectâ← s material settings. The key light gives that â← in-the-spotlightâ← feel, highlighting the subject, while the backlight gives a crisp edge to the subject against the background. This helps them stand out.

In this rig, the key light is a fairly bright spot light with an energy of **2.0** and a falloff Dist of **10**. Use a slighter tinge of yellow because the light is so bright; it is the only light for that side. The other sidelight has been moved in back and raised to eye (camera) level. You need to cut the energy of the backlight in half, or when it is added to the remaining sidelight, it will light up the side too much and call too much attention to itself. You can vary the angle and height of the backlight to mimic a sun lighting up the objects.

Use this rig in normal 3D animations to light the main actor. Use this rig especially if you have transparent objects (like glass) so that there is plenty of light to shine through them to the camera. The tricky part here is balancing the intensities of the lights so that no one light competes with or overpowers the others, while making sure all three work together as a team.

Four-point Rig



Four-point light rig.
The four-point lighting rig provides a better simulation of outside lighting, by adding a Sun lamp **30** Blender Units above, **10** to the side, and **15 BU** behind the subject. This sunlight provides backlighting and fills the top of the subject; even producing an intentional glare on the top of their head, telling you there is a sun up there. Notice it is colored yellow, which balances out the blue sidelights.

Changing the key light to a Spot Inverse Square No Specular, with an energy of **1.0** and pure white light that has a falloff Dist of **12.0** combines with and softens the top sun flare while illuminating the face, resulting in a bright sunshine effect. Two lights above means sharper shadows as well, so you might want to adjust the side fill lights. In this picture, they are still Hemi, No Specular with an energy of **0.1**, falloff Distance of **10.0**, and blueish speckled as before.

Use this rig when the camera will be filming from behind the characters, looking over their shoulder or whatnot, because the sun provides the backlight there. Also use this rig when you have transparent objects, so there is light to come through the objects to the camera.

Another spot for the fill light is shining up onto the main actorâ← s face, illuminating the underside of his chin and neck. This gets rid of a sometimes ugly shadow under the chin, which if not corrected, can make the actor look fat or like they have a double chin; otherwise distracting. It evens out the lighting of the face.

Troubleshooting and a Helpful Blend file

If you run into a problem with your render, where there are really bright areas, or really dark ones, or strange shadows, or lines on your objects, here is what I suggest you do:

1. First, try deactivating all materials (create a default, gray one, and enter its name in the Mat field, Render Layers panel, Render sub-context â← Â to get back all your normal materials, just erase this text field!). See if you get those problems with just grayness objects. If you donâ← t have the problem anymore, that should tell you that youâ← ve got a materials-interacting-with-light problem. Check the material settings, especially ambient, reflection and all those little buttons and sliders on the Shaders panel. You can set some lights to affect only certain materials, so if thereâ← s an issue with only a few objects being really bright, start with those.
2. Then start â← killingâ← lights (e.g. moving them to an unused layer); regress all the way back to one light, make sure itâ← s smooth, then add them in one by one. As they add together, reduce power in the tested ones so they merge cleanly, or consider not adding it at all, or, especially, reduce the energy of the lamp you just introduced.
3. You can also set lights to only light objects on a layer, so again, if some of the gray spheres have weirdness, check for that as well. Again, you may have done some of this accidentally, so sometimes deleting the light and re-adding it with defaults helps you reset to a known-good situation.
4. Negative lights can be very tricky, and make your model blotchy, so pay special attention to your use of those special lights. Shadow-only lights can throw off the look of the scene as well. Overly textured lights can make your scene have random weird colors. Donâ← t go too far off a slight tinge of blue or yellow or shades of white, or your material may show blue in the Material sub-context but render green, and you will be very confused.
5. Look at your environment settings (World sub-context, World panel): Horizon, Zenith, and Ambient light.

All of the above lighting rigs were done without raytracing or using Yafray or nodes; those are additional complication layers that I cannot cover here. Sorry.

Hopefully, you will want to download the blend file here. Save it on your hard drive, and when you want to start a new scene, do a File Â» Append Â» `(filename)` Â» Scene Â» `(and select the rig you want)`.

Introduction



Radiosity example.
Most rendering models, including ray-tracing, assume a simplified spatial model, highly optimized for the light that enters the camera â← eyeâ← in order to draw the image. You can add reflection and shadows to this model to achieve a more realistic result. Still, thereâ← s an important aspect missing! When a surface has a reflective light component, it not only shows up in our image, it also shines light at surfaces in its neighborhood. And *vice-versa*. In fact, light bounces around in an environment until all light energy is absorbed (or has escaped!). Re-irradiated light carries information about the object which has re-irradiated it, notably color. Hence not only the shadows are â← less blackâ← because of re-irradiated light, but also they tend to show the color of the nearest, brightly illuminated, object. A phenomenon often referred to as â← color leakingâ← (*Radiosity example*).

In closed environments, light energy is generated by â← emittersâ← and is accounted for by reflection or absorption of the surfaces in the environment. The rate at which energy leaves a surface is called the â← radiosityâ← of a surface. Unlike conventional rendering methods, radiosity methods first calculate all light interactions in an environment in a view-independent way. Then, different views can be rendered in real-time. In Blender, since version 2.28, radiosity is both a rendering and a modeling tool. This means that you can enable radiosity within a rendering or rather use radiosity to paint vertex colors and vertex lights of your meshes, for later use.

Normal directions are crucial
The light emission is governed by the direction of the normals of a mesh. This is of course crucial for emitters, but also for all meshes, as when they re-irradiate their unshot energy, they do it in their normal direction.

Technical Details

During the late eighties and early nineties, radiosity was a hot topic in 3D computer graphics. Many different methods were developed, the most successful of these solutions were based on the â← progressive refinementâ← method with an â← adaptive subdivisionâ← scheme. And this is what Blender uses. To be able to get the most out of the Blender radiosity method, it is important to understand the following principles:

Finite Element Method

> Many computer graphics or simulation methods assume a simplification of reality with â← finite elementsâ← . For a visually attractive (and even scientifically proved) solution, it is not always necessary to dive into a molecular level of detail. Instead, you can reduce your problem to a finite number of representative and well-described elements. It is a common fact that such systems quickly

converge into a stable and reliable solution. The radiosity method is a typical example of a finite element method inasmuch as every face is considered a â← finite elementâ← and its light emission considered as a whole.

Patches and Elements

In the radiosity universe, we distinguish between two types of 3D faces:

Patches

These are triangles or squares which are able to *send energy*. For a fast solution it is important to have as few of these patches as possible. But, to speed things up the energy is modeled as if it were radiated by the patchâ← s center; the size of the patches should then be small enough to make this a realistic energy distribution. (For example, when a small object is located above the patch center, all energy the patch sends is obscured by this object, even if the patch is larger! This patch should be subdivided in smaller patches).

Elements

These are the triangles or squares which *receive energy*. Each Element is associated with a Patch. In fact, patches are subdivided into many small elements. When an element receives energy it absorbs part of it (depending on its color) and passes the remainder to the patch, for further radiation. Since the elements are also the faces that we display, it is important to have them as small as possible, to express subtle shadow boundaries and light gradients.

Progressive Refinement

This method starts with examining all available patches. The one with the most â← unshotâ← energy is selected to shoot all its energy to the environment. The elements in the environment receive this energy, and add part of it to the â← unshotâ← energy of their associated patches (the remaining being absorbed). Then the process starts again for the patch *now* having the most unshot energy. This continues for all the patches until no energy is received anymore, or until the â← unshotâ← energy has converged below a certain value.

The hemicube method

The calculation of how much energy each patch gives to an element is done through the use of â← hemicubesâ← . Exactly located at the patchâ← s center, a hemicube (literally, â← half a cubeâ←) consist of five small images of the environment. For each pixel in these images, a certain visible element is color-coded, and the transmitted amount of energy can be calculated. Especially with the use of specialized hardware, the hemicube method can be accelerated significantly. In Blender, however, hemicube calculations are done â← in softwareâ← . This method is in fact a simplification and optimization of the â← realâ← radiosity formula (form factor differentiation). For this reason, the resolution of the hemicube (the number of pixels of its images) is approximated, and its careful setting is important to prevent aliasing artifacts.

Adaptive subdivision

Since the size of the Patches and Elements in a mesh defines the quality of the radiosity solution, automatic subdivision schemes have been developed to define the optimal size of patches and elements. Blender has two automatic subdivision methods:

1. Subdivide-shoot Patches

By shooting energy to the environment, and comparing the hemicube values with the actual mathematical â← form factorâ← value, errors can be detected that indicate a need for further subdivision of the patch. The results are smaller patches and a longer solving time, but a higher realism of the solution.

2. Subdivide-shoot Elements

By shooting energy to the environment, and detecting high energy changes (gradients) inside a patch, the elements of this patch are subdivided one extra level. The results are smaller elements and a longer solving time and maybe more aliasing, but a higher level of detail.

Display and Post Processing

Subdividing elements in Blender is â← balancedâ← , that means each element differs a maximum of *one* subdivide level with its neighbors. This is important for a pleasant and correct display of the radiosity solution with Gouraud shaded faces. Usually, after solving, the solution consists of thousands of small elements. By filtering these and removing â← doublesâ← , the number of elements can be reduced significantly without destroying the quality of the radiosity solution. Blender stores the energy values in â← floating pointâ← values. This makes settings for dramatic lighting situations possible, by changing the standard multiplying and gamma values (this is the same principle as HDR with modern GPUs).

Radiosity for Modeling

The final step can be replacing the input meshes with the radiosity solution (button Replace Meshes, Radio Tool panel, Radiosity sub-context, Shading context, F5). At that moment the vertex colors are converted from a â← floating pointâ← value to a 24 bits RGB value. The old mesh objects are deleted and replaced with one or more new mesh objects. You can then delete the radiosity data with Free Data. The new objects get a default material that allows immediate rendering. Two settings in a material are important for working with vertex colors (both in Material panel, Material sub-context):

VCol Paint

This option treats vertex colors as a replacement for the normal RGB value in the material. You have to add lamps in order to see the radiosity colors. In fact, you can use Blender lighting and shadowing as usual, and still have a neat radiosity â← lookâ← in the rendering.

VCol Light

The vertex colors are added to the light when rendering. Even without Lamps, you can see the result. With this option, the vertex colors are pre-multiplied by the material RGB color. This allows fine-tuning of the amount of â← radiosity lightâ← in the final rendering.

Radiosity Rendering

Mode: All modes

Panel: Render (Scene context, Render sub-context, F10)
Radio Render (Shading context, Radiosity sub-context, F5)

Description

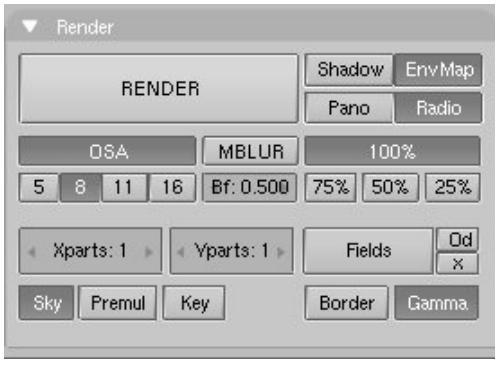
Letâ← s assume you have a scene ready, and that you want to render it with radiosity. The first thing to grasp when doing radiosity is that *no lamps are necessary*, but some meshes with an Emit material property greater than zero are required, since these will be the light sources. Emit is found in the bottom right of the Shaders panel (Material sub-context). Typically, a value of **0.5** or less gives a soft radiance.

You can build the test scene shown in (*Set-up for radiosity test*), it is rather easy. Just make a big cube for the room (with *switched normals*, i.e. pointing inwards!), give different materials to the side walls, add a cube and a stretched cube within it, and add a plane with a non-zero Emit value next to the roof, to simulate the area light (with its normal pointing down).

You assign materials as usual to the input models. The RGB value of the material defines the patch color. The Emit value of a material defines if a patch is loaded with energy at the start of the radiosity simulation. The Emit value is multiplied with the area of a patch to calculate the initial amount of unshot energy â← Â see the technical details section in the preceding page.

Emitting faces
Check the number of â← emittersâ← on Blender console! If this is zero nothing interesting can happen. You need at least one emitting patch to have light and hence a solution.

Enabling Radiosity



The Render panel with Radio button highlighted.
The Radio button on Render panel (Scene context, Render sub-context) enables radiosity calculations as part of the render process. This will automatically consider all objects in the scene, and those materials with an non-null Emit value will emit light onto other objects.

Material Options

When assigning materials, be sure that all of them have Radio enabled on the Links and Pipeline panel (it is set by default).

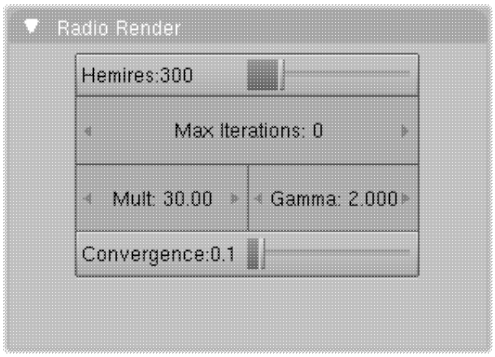
You also need to have the Amb setting of each material above zero, for it to receive emitted light, since the emitted light becomes part of the ambient light. Ambient is found on the Shaders panel.

Objects that emit, or radiate light, should have non-null Emit as well, also found on the Shaders panel.

Since all objects realistically reflect some light back into the environment, a good general practice, for every objectâ← s material, is to always:

- Enable Radio.
- Set Ambient > 0, typically **0.1**.
- Set Emit > 0, typically **0.1**.

Rendering Options



The Radio Render panel.
The radiosity specific options for render-time radiosity are all found in the Radio Render panel (Shading context, Radiosity sub-context, F5). The other Radio Tool panel is only relevant when using radiosity as a modeling tool (see next page).

Hemires
: The hemicube resolution; the color-coded images used to find the elements that are visible from a â← shoot patchâ← , and thus receive energy. Hemicubes are not stored, but are recalculated each time for every patch that shoots energy. The Hemires value determines the radiosity quality and adds significantly to the solving time.

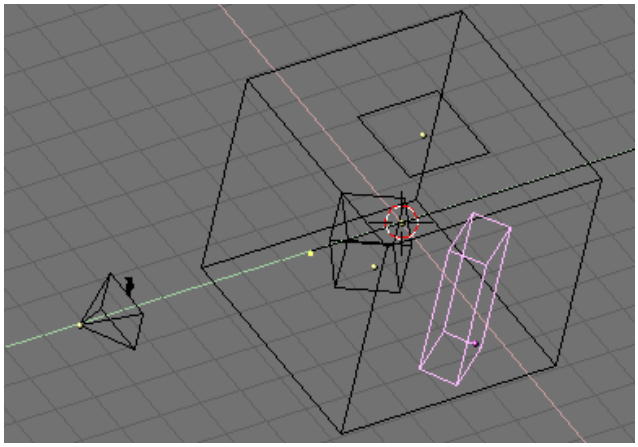
Max Iterations
: The maximum number of radiosity iterations. If set to zero (the default), radiosity will go on until the convergence criterion is met. You are strongly advised to set this to some non-zero number, usually greater than **100**.

Mult, Gamma
: The color space of the radiosity solution is far more detailed than can be expressed with simple 24 bit RGB values. When elements are converted to faces, their energy values are converted to an RGB color using the Mult and Gamma values. With the Mult value you can multiply the energy value (i.e. increase/decrease the luminosity of the whole picture), with Gamma you can change the contrast of

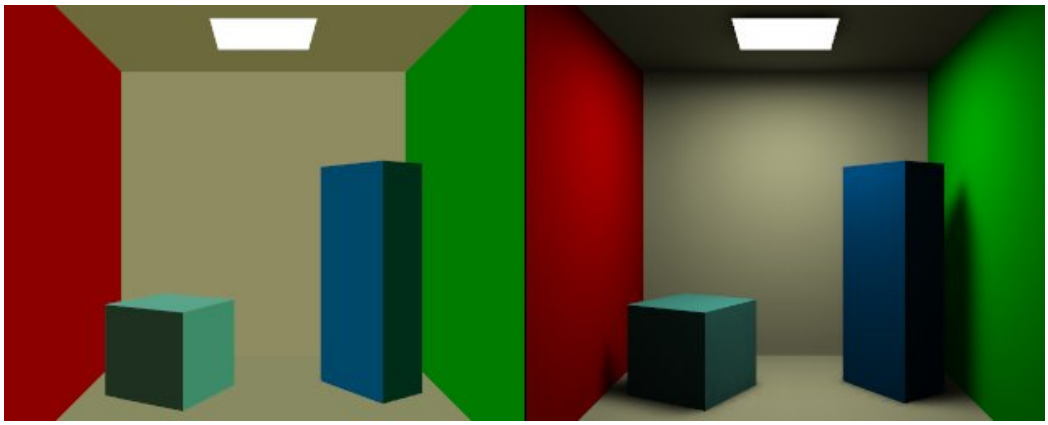
the energy values (lower values mean higher contrast).

Convergence
: When the amount of unshot energy in an environment is lower than this value, the radiosity solving stops. The initial unshot energy in an environment is multiplied by the area of the patches. During each iteration, some of the energy is absorbed, or disappears when the environment is not a closed volume. In Blenderâ← s standard coordinate system a typical emitter (as in the examples) has a relatively small area. The Convergence value is divided by a factor of **1000** before testing for that reason.

Examples



Set-up for radiosity test.
The rendering will take more time than usual, in the console you will notice a counter going up (as well as the mouse pointer). The result will be quite poor (*Radiosity rendering for coarse meshes (left) and fine meshes (right)*, left), because the automatic radiosity render does not do adaptive refinement! Select all meshes, one after the other, and in Edit mode, subdivide them at least three times â← Â or add to them Subsurf modifiersâ←¦ The room, which is much bigger than the other meshes, you can even subdivide four times. Set the Max Iterations a bit higher, **300** or more. Try rendering again (F12). This time the rendering will take even longer, but the results will be much nicer, with soft shadows and color leakage (*Radiosity rendering for coarse meshes (left) and fine meshes (right)*, right).



Radiosity rendering for coarse meshes (left) and fine meshes (right).

Note
In the radiosity rendering, Blender acts as for a normal rendering, this means that textures, curves, surfaces

and even dupliframed objects are handled correctly.

Hints

Recall that **the light emission is governed by the direction of the normals of a mesh**. This is of course crucial for emitters, but also for all meshes, as when they re-irradiate their unshot energy, they do it in their normal direction â← Â just try to render the above example with the normal of the â← lightâ← plane pointing upward, or the normals of the â← roomâ← cube pointing outwardsâ←¦

Itâ← s also very important to understand that render radiosity does no automatic subdividing to get patches and elements of a reasonably good size. You have to do it by hand, especially when using simple objects with very large faces (like the â← roomâ← cube in example above), either by subdividing in Edit mode, or by adding Subsurf modifiersâ←¦

As with everything in Blender, Radiosity settings are stored in a datablock. It is attached to a scene, and each scene in Blender can have a different radiosity â← blockâ← . Use this facility to divide complex environments into scenes with independent radiosity solvers.

Baking Radiosity

Radiosity Tool

Mode: All modes

Panel: Radio Render, Radio Tool, and Calculation (Shading context, Radiosity sub-context, F5)

Description

Radiosity can be used also as a tool for defining vertex colors and lights. This can be very useful if you want to make further tweaks to your models, or use them in the Game Engine. Furthermore, the radiosity modeling allows for adaptive refinement, whereas the radiosity rendering does not! There are few important points to grasp for practical radiosity modeling: Only meshes in Blender are allowed as input for radiosity modeling. This because the process generates vertex colorsâ←¦ and so there must be vertices. It is also important to realize that *each* face in a mesh becomes a patch, and thus a potential energy emitter and reflector. Typically, large patches send and receive more energy than small ones. It is therefore important to have a well-balanced input model with patches large enough to make a difference! When you add extremely small faces, these will (almost) never receive enough energy to be noticed by the â← progressive refinementâ← method, which only selects patches with large amounts of unshot energy.

Non-mesh objects
Only meshes means that you have to convert curves and surfaces to meshes before starting the radiosity solution!

Collecting Meshes

Mode: All modes

Panel: Radio Tool (Shading context, Radiosity sub-context, F5)

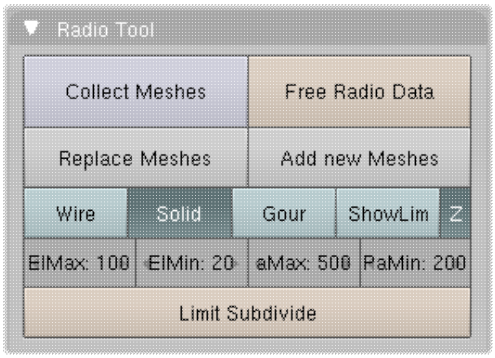
Description

The first step in the process is to convert the selected meshes to radiosity patches, to participate in the solution.

Options



Radio Tool panel, before meshes collecting.



Radio Tool panel, after meshes collecting.

Collect Meshes
: Convert all selected and visible meshes in the current scene to patches. As a result, Collect Meshes is no more highlighted, and a new panel, Calculation, appears. Blender has now entered the Radiosity Modeling mode, and other editing functions are blocked until the newly highlighted button Free Data has been pressed.

After the meshes are collected, they are drawn in one of the three pseudo lighting modes, that clearly differ from the normal drawing. The Radio Tool panel has three buttons controlling this:

Wire, Solid, Gouraud
: These control the radiosity modeling draw mode to be used, independently from the indicated draw mode of a 3D window. Wire gives a wire view of the meshes, with the subdivided elements produced during calculations, and colored according to the computed solution. Solid shows all elements as solid flat faces, without blending between neighbor elementsâ← colors. Gouraud display is only performed after the radiosity process has started â← Â use it to have nice, smooth colored surfaces.

Subdivision limits

Mode: All modes

Panel: Radio Tool (Shading context, Radiosity sub-context, F5)

Description

Blender offers a few settings to define the minimum and maximum sizes of patches and elements.

Options



Radio Tool panel with subdivision controls highlighted.

Limit Subdivide
: With respect to the values PaMax and PaMin, the patches are subdivided. This subdivision is also automatically performed when a GO action has started (Calculation panel).

ElMax, ElMin, PaMax, PaMin
: The maximum and minimum size of elements and patches. These limits are used during all radiosity phases. The unit is expressed in **0.0001** of the bounding box size of the entire environment. Hence, with default **500** and **200** maximum and minimum settings, patch size is between **0.05** (1/20th of the entire model) and **0.02** (1/50th of the entire model).

ShowLim, Z
: These buttons allow you to visualize the patch and element limits. By pressing ShowLim only, you get the limits drawn in the global XY plane. Z only gives limits drawn in the global XZ plane, and with both enabled, in the global YZ plane. The white lines show the patch limits, cyan lines show the element limits.

Adaptive Subdivision

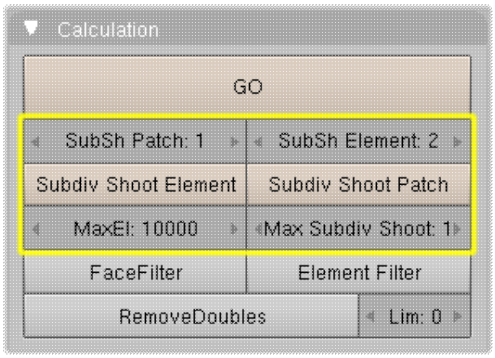
Mode: All modes

Panel: Calculation (Shading context, Radiosity sub-context, F5)

Description

The Calculation panel regroups the last settings before starting the calculations â← Â those controlling the adaptive subdivision of elements and patches.

Options



The Calculation panel, with adaptive subdivision settings highlighted.

MaxEl
: The maximum allowed number of elements, from **1** to **250Â 000**. Since elements are subdivided automatically in Blender, the amount of used memory and the duration of the solving time can be controlled with this button. As a rule of thumb, **20Â 000** elements take up **10 Mb** memory.

Max Subdiv Shoot
: The maximum number of shoot patches that are evaluated for the â← adaptive subdivisionâ← , from **1** to **250**.

Subdiv Shoot Patch
: By shooting energy to the environment, errors can be detected that indicate a need for further subdivision of patches. The subdivision is performed only once each time you call this function. The results are smaller patches and a longer solving time, but a higher realism of the solution. This option will also be automatically performed when the GO action has started.

Subdiv Shoot Element
: By shooting energy to the environment, and detecting high energy changes (frequencies) inside a patch, the elements of this patch are selected to be subdivided one extra level. The subdivision is performed only once each time you call this function. The results are smaller elements and a longer solving time and probably more aliasing, but a higher level of detail. This option will also be automatically performed when the GO action has started.

SubSh Patch
: The number of times the environment is tested to detect patches that need subdivision.

SubSh Element
: The number of times the environment is tested to detect elements that need subdivision.

Computing Radiosity

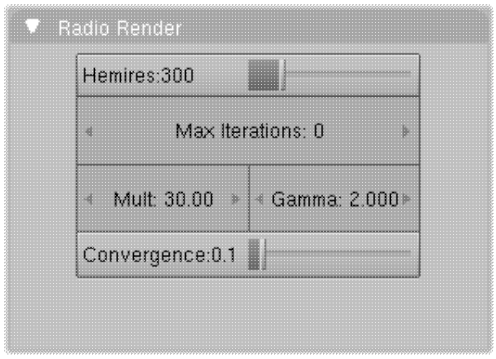
Mode: All modes

Panel: Radio Render, Calculation (Shading context, Radiosity sub-context, F5)

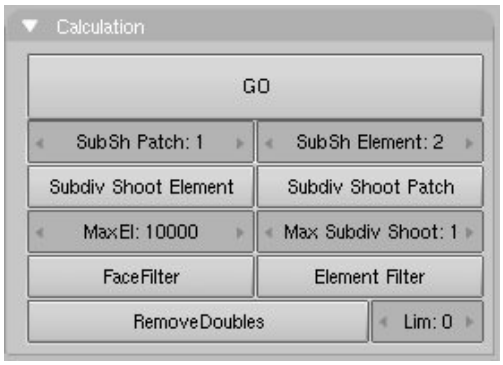
Description

This is where the radiosity solution is actually computed. Youâ← ll find here the same options as those for the direct radiosity rendering.

Options



The Radio Render panel.



The Calculation panel.

Hemires, Max Iteration, Convergence
: The hemicube resolution, maximum number of radiosity iterations, and maximum remaining level of unshot energy. These three settings have already been described in the previous page, and are still active here, with the same meaning.

GO
: Start the radiosity simulation. The phases are:
 - ◊ *Limit Subdivide*. When patches are too large, they are subdivided.
 - ◊ *Subdiv Shoot Patch*. The value of SubSh Patch defines the number of times the Subdiv Shoot Patch function is called. As a result, patches are subdivided.
 - ◊ *Subdiv Shoot Element*. The value of SubSh Element defines the number of times the Subdiv Shoot Element function is called. As a result, elements are subdivided.
 - ◊ *Subdivide Elements*. When elements are still larger than the minimum size, they are subdivided. Now, the maximum amount of memory is usually allocated.
 - ◊ *Solve*. This is the actual â← progressive refinementâ← method. The mouse pointer displays the iteration step, the current total of patches that shot their energy in the environment. This process continues until the unshot energy in the environment is lower than the Convergence value, or when the maximum number of iterations has been reached.
 - ◊ *Convert to faces*. The elements are converted to triangles or squares with â← anchoredâ← edges, to make sure a pleasant not-discontinue Gouraud display is possible.

 This process can be terminated with Esc during any phase.

Editing the solution

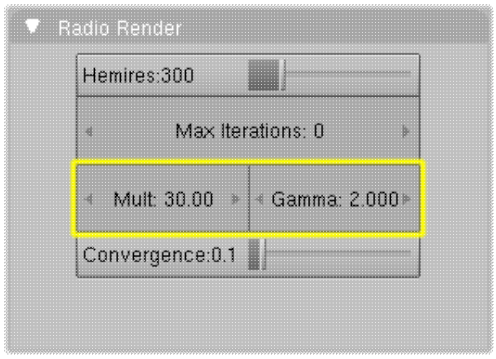
Mode: All modes

Panel: Radio Render, Radio Tool, Calculation (Shading context, Radiosity sub-context, F5)

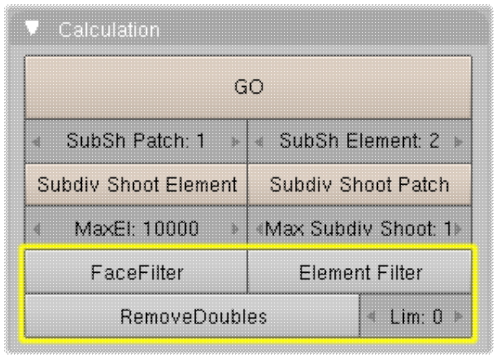
Description

Once the radiosity solution has been computed, there are still some actions to take.

Options



The Radio Render panel.



The Calculation panel.



The Radio Tool panel.

The radiosity panels, with post-processing settings highlighted.

Mult, Gamma
: These two options from the Radio Render panel control the luminosity and contrast of the result, as described in previous page.

Other post-processing features, specific to radiosity modeling, are found in the Calculation panel:

Element Filter
: This option filters elements to remove aliasing artifacts, to smooth shadow boundaries, or to force equalized colors for the RemoveDoubles option.

RemoveDoubles, Lim
: When two neighboring elements have a displayed color that differs less than the limit specified in the Lim setting, the elements are joined. The Lim value ranges from **0** (same colors) to **50**.

FaceFilter
: Elements are converted to faces for display. A FaceFilter forces an extra smoothing in the displayed result, without changing the element values themselves.

Finally, you must â← applyâ← your radiosity modeling to your scene, in the Radio Tool panel:

Add New Meshes
: The faces of the current displayed radiosity solution are converted to mesh objects with vertex colors. A new material is added that allows immediate rendering. *The input-meshes remain unchanged.*

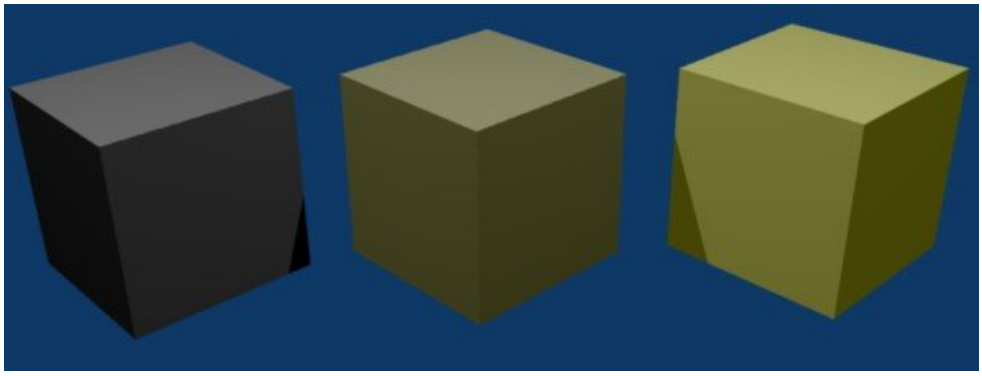
Replace Meshes
: As previous, but *the input-meshes are removed/deleted*.

Free Radio Data
: All patches, elements and faces are freed in memory. You must always perform this action after using radiosity to be able to return to normal editing.

Ambient Light

Description

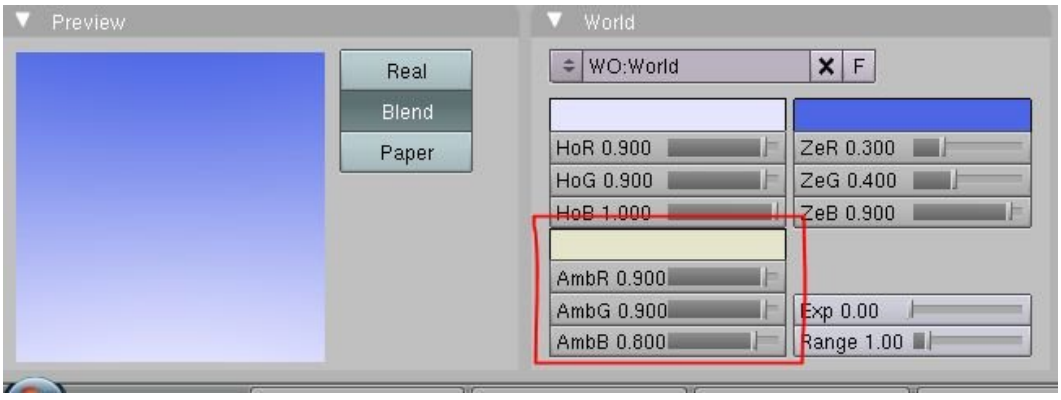
Ambient light is all around us and is the result of the sun and lamps scattering photons every which way, reflecting off of and wavelengths being absorbed by objects. Rather than try to calculate the exact intensity of each and every photon, use the Ambient light settings to generally light the scene.



Ambient settings of **0.00**, **0.25**, and **0.50**, receiving a yellow light with RGB values of (**0.5**,**0.5**,**0.0**).

You can specify the amount of Ambient light (the color of which is set in the world settings); use an off-white to simulate the color of the ambient lighting inside closed rooms or on different planets based on its sunâ← s visible spectrum. Generally, about half of RGB gives a nice soft lighting to a scene.

Setting Ambient Light



Suggested Earth sunny day settings.

You set the color of ambient light by clicking on the color swatch and using the color picker to choose your color, or by manually adjusting the RGB sliders, or LMB clicking on the slider value (number) and entering a number using your keyboard. These fields are outlined in red in the example above (World panel, World sub-context).

This simple example of ambient light settings mimics Earth on a sunny day. The color of the sky at the horizon is white with a touch of haze and a tinge of blue, and the zenith overhead is a dark blue. The sun is out and bright, with a yellow tinge, as on a cloudless day. Think Caribbean, and youâ← ll get the idea. London in the fog is grey; night-time is black, etc.

Material Ambient Effect



The Shaders panel with Ambient slider highlighted.
Each objectâ← s material settings (Shaders panel) has an Ambient slider that lets you choose how much ambient light that object receives. Generally, about **0.1** to half is good, depending on the color of the ambient light. A setting of **0.0** means that the object does not receive any ambient light; it is only lit by light that actually gets to it. A setting of **1.0** means that it is lit by all of the worldâ← s ambient light. Increasing the Ambient setting has the effect of flattening the shading, since the ambient light washes out the diffuse shading.

You should set the material slider as shown above in the three cube example, based on the amount of ambient light you think the object will receive. Something deep in the cave will not get any ambient light, whereas something close to the entrance will get more. Note that you can animate this effect, to change it as the object comes out of the shadows and into the light.

World Ambient Occlusion

A trick for games is to illuminate surfaces based on how far they are away from occluding geometry, called ambient occlusion. If looks like ambient light is lighting the objects in the scene, and is the process of Ambient Occlusion. The color of the â← lightâ← applied to surfaces is affected by the world ambient color. See wiki page on Ambient Occlusion.

Radiosity

Radiosity is the process of applying radiated light from objects in the scene as they add to the color of nearby ambient light, which in turn affects the color of the ambient light. See wiki page on Radiosity.

Ambient Occlusion

Mode: All modes

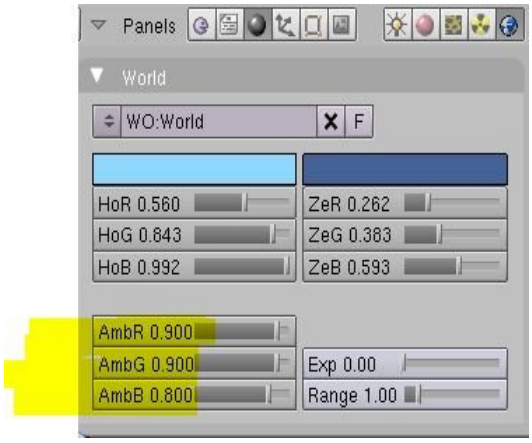
Panel: Amb Occ (Shading Context, World sub-context, F8)

Description

Ambient Occlusion is a sophisticated raytracing calculation which simulates soft global illumination by faking darkness perceived in corners and at mesh intersections, creases, and cracks, where light is diffused (usually) by accumulated dirt and dust. This has the effect of darkening cracks, corners and points of contact, which is why ambient occlusion is often referred to as a â← dirt shaderâ← .

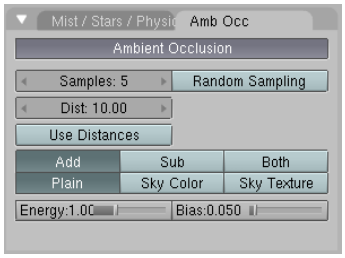
There is no such thing as AO in real life, AO is a specific not-physically-accurate (but generally nice looking) rendering trick. It basically samples a hemisphere around each point on the face, sees what proportion of that hemisphere is occluded by other geometry, and shades the pixel accordingly. Itâ← s got nothing to do with light at all, itâ← s purely a rendering trick that tends to look nice because generally in real life surfaces that are close together (like small cracks) will be darker than surfaces that donâ← t have anything in front of them, because of shadows, dirt, etc. The AO process though is approximating this result, itâ← s not simulating light bouncing around or going through things. Thatâ← s why AO still works when you donâ← t have any lights in the scene, and itâ← s why just switching on AO alone is a very bad way of â← lightingâ← a scene.

You must have raytracing enabled as a Render panel option (Scene context, F10, Render sub-context) for this to work.



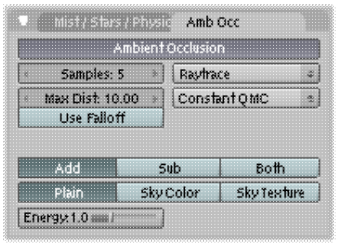
The World panel with ambient color sliders highlighted.
You must have an ambient light color (Shading context, World sub-context, F8, AmbRGB sliders) set to your desires. By default, the ambient light color (world) is black, simulating midnight in the basement during a power outage. Applying that color as ambient will actually darken all colors. A good outdoor mid-day color is RGB (**0.9**, **0.9**, **0.8**) which is a whitish yellow sunny kind of color on a bright-but-not-harshly-bright day.

Options



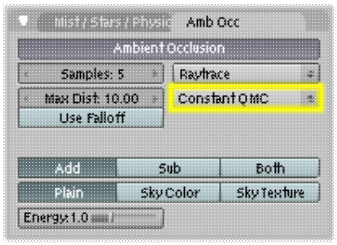
The Amb Occ panel, method selection.
Ambient occlusion has evolved a bit in Blender 2.49. You have now two main methods: Raytrace and Approximate. For those who knew the tool in previous versions (2.48 and below), they will find nearly the same options (with the same behavior) using the Raytrace method, Constant Jittered sample generator.

Raytrace



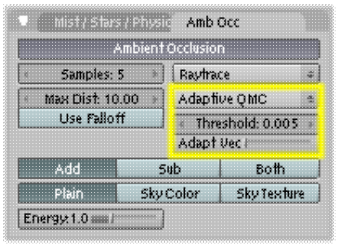
The Amb Occ panel, Raytrace method.
The Raytrace method gives the more accurate, but also the more noisy results. You can get a nearly noiseless image, but at the cost of render timeâ←¦ It is the only option if you want to use the colors of your skyâ← s texture.

You have the three standard sample options:



Constant QMC sampling.

Constant QMC
: The base Quasi-Monte Carlo, gives evenly and randomly distributed rays.



Adaptive QMC sampling.

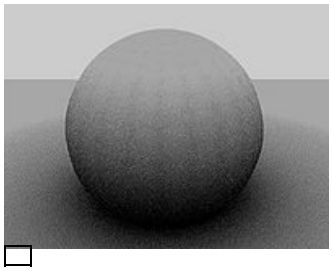
Adaptive QMC
: An improved version of QMC, that tries to determine when sample can be lowered or skipped, based on its two settings:
 Threshold
 : The limit below which the sample is considered fully occluded (â← blackâ←) or un-occluded (â← whiteâ←), and skipped.
 Adapt Vec
 : A factor to reduce AO sampling on fast moving pixels. As it uses the Vec render pass, this one must be enabled (see render passes page).

About QMC
: See also the raytraced shadows page for more info about Qausi-Monte Carlo sampling method.

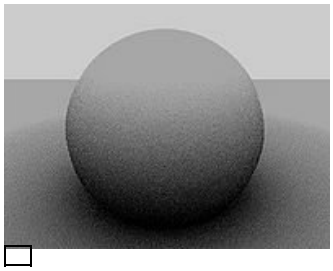
Constant Jittered
: The historical sample method, more prone to â← biasâ← artifactsâ←¦
 Bias
 : The angle (in radians) the hemisphere will be made narrower (i.e. the hemisphere will be no more a real hemisphere: its section will be no more a semicircle, but an arc of a circle of â← *pi* - Biasâ← radians).
 The bias setting allows you to control how smooth â← smoothâ← faces will appear in AO rendering. Since AO occurs on the original faceted mesh, it is possible that the AO light makes faces visible even on objects with â← smoothâ← on. This is due to the way AO rays are shot, and can be controlled with the Bias slider. Note that even if it might happens with QMC sampling methods, this is much more visible with the Constant Jittered one â← Â and anyway, you have no Bias option for QMCâ←¦



Constant Jittered sampling.

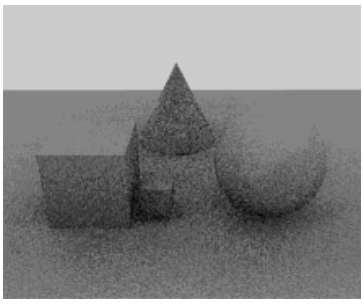


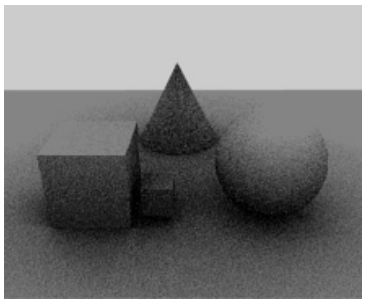
24Ã 24 UV Sphere with Bias: **0.05** (default). Note the facets on the sphereâ← s surface even though it is set smooth.

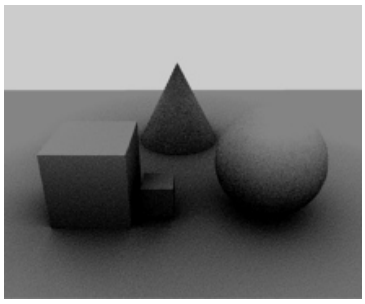


Raising the Bias to **0.15** removes the faceted artifacts.

Samples
: The number of rays used to detect if an object is occluded. Higher numbers of samples give smoother and more accurate results, at the expense of slower render times. The default value of **5** is usually good for previews. The actual amount of shot rays is the square of this number (i.e. Samples at **5** means **25** rays). Rays are shot at the hemisphere according to a random pattern (determined by the sample methods described above), this causes differences in the occlusion pattern of neighboring pixels unless the number of shot rays is big enough to produce good statistical data.

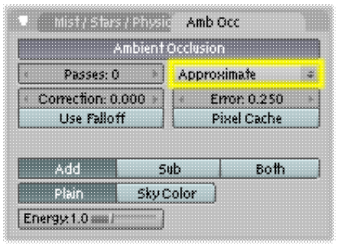

Ambient Occlusion with **3** Samples.


Ambient Occlusion with **6** Samples.


Ambient Occlusion with **12** Samples.

Max Dist (formerly Dist)
: The length of the occlusion rays. The longer this distance, the greater impact that far away geometry will have on the occlusion effect. A high Max Dist value also means that the renderer has to search a greater area for geometry that occludes, so render time can be optimized by making this distance as short as possible, for the visual effect that you want.

Approximate



The Amb Occ panel, Approximate method.
The Approximate method gives a much smoother result for the same amount of render time, but as its name states, it is only an approximation of the Raytrace method, which implies it might produce some artifacts â← Â and it cannot use the skyâ← s texture as base colorsâ←¦

I donâ← t know how it works precisely, so if anybody does, please write here a small description of the algorithmâ←¦

Error
: This is the tolerance factor for approximation error (i.e. the max allowed difference between approximated result and fully computed result). The lower, the slower the render, but the more accurate the resultsâ←¦ Ranges between **0.0** and **10.0**, defaults to **0.250**.

Pixel Cache
: When enabled, it will keep values of computed pixels, to interpolate it in its neighbors. This speeds up further more the render, generally without visible loss in qualityâ←¦

This method seems to tend to â← over-occludeâ← the results. You have two complementary options to reduce this problem:

Passes
: Set the number of pre-processing passes, between **0** (no pre-processing) to **10**. Donâ← t know what this pre-processing does, thoughâ←¦

Correction
: A correction factor to reduce over-occlusion. Ranges between **0.0** (no correction) to **1.0**.

Common Settings

Use Falloff (formerly Use Distance
: When activated, the distance to the occluding objects will influence the â← deepnessâ← of the shadow. This means that the farer away is an occluding geometry, the lighter will be its â← shadowâ← . This effect only occurs when the Strength factor is higher than **0.0**. It mimics the light dispersion in the atmosphereâ←¦

Stength (formerly DistF)
: Controls the attenuation of the shadows enabled with Use Falloff. Higher values give a shorter shadow, as it falls off more quickly (corresponding to a more foggy/dusty atmosphere). Ranges from **0.0** (default, no falloff) to **10.0**.

Blending

The ambient occlusion pass is composited during the render pipeline. Three blending modes are available:

Add
: The pixel receives light according to the number of non-obstructed rays. The scene is lighter.

Sub
: The pixel receives shadow (negative light) according to the number of obstructed rays. The scene is darker.

Both
: Both effects take place, the scene has more or less the same brightness, but more contrast.

Energy
: The strength of the AO effect, a multiplier for addition or subtraction.

Note
If Sub is chosen, there must be other light sources, otherwise the scene will be pitch black. In the other two cases the scene is lit even if no explicit light is present, just from the AO effect. Although many people like to use AO alone as a quick shortcut to light a scene, the results it gives will be muted and flat, like an overcast day. In most cases, it is best to light a scene properly with Blenderâ← s standard lamps, then use AO on top of that, set to â← Subâ← , for the additional details and contact shadows.

Ambient Color

Ambient occlusion can take the color of its lighting from various sources:

Plain
: The pixel receives shading based on the worldâ← s ambient color.

Sky Color

The pixel receives shading based on the worldâ← s sky color. The color is computed on the basis of the portion of the sky hit by the non-obstructed rays (*Ambient occlusion with Sky Color. Zenith is blue, Horizon is orange, and type is Blend so that sky goes full orange at nadir.*).

Sky Texture

A sky image texture must be present, possibly an AngMap or a SphereMap. It behaves as Sky Color, but the ray color depends on the color of the sky texture pixel hit (only available with the Raytrace method).



Ambient occlusion with Sky Color. Zenith is blue, Horizon is orange, and type is Blend so that sky goes full orange at nadir.



Ambient occlusion with Sky Texture, using a scaled down St. Peters Basilica HDR AngMap [1].

Technical Details

Ambient occlusion is calculated by casting rays from each visible point, and by counting how many of them actually reach the sky, and how many, on the other hand, are obstructed by objects.

The amount of light on the point is then proportional to the number of rays which have â← escapedâ← and have reached the sky. This is done by firing a hemisphere of shadow-rays around. If a ray hits another face (it is occluded) then that ray is considered â← shadowâ← , otherwise it is considered â← lightâ← . The ratio between â← shadowâ← and â← lightâ← rays defines how bright a given pixel is.

Hints

Ambient occlusion is a raytracing technique (at least with Raytrace method), so it tends to be slow. Furthermore, performance severely depends on octree size, see the rendering chapter for more information.

Exposure and Range

Mode: All modes

Panel: World (Shading context, World sub-context, F8)

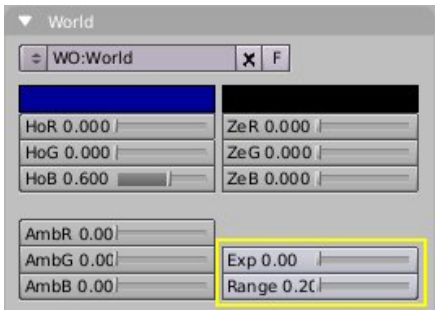
Description

Exposure and Range are similar to the â← Color Curvesâ← tool in Gimp or Photoshop.

Previously Blender clipped color straight with â← **1.0**â← (or 255) when it exceeded the possible RGB space. This caused ugly banding and overblown highlights when light overflowed (*An overexposed teapot*).

Using an exponential correction formula, this now can be nicely corrected.

Options



Exposure and Range sliders.

Exp
: The exponential curvature, with **0.0** being linear, and **1.0** being curved.

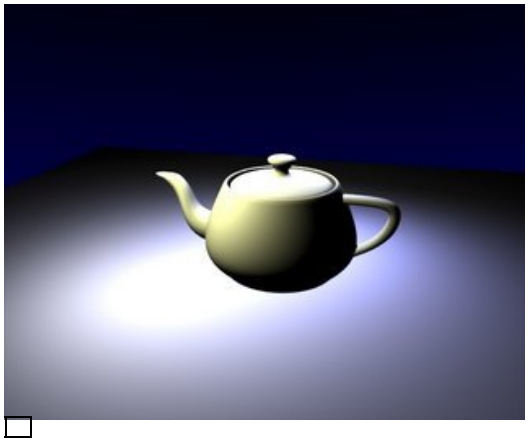
Range
: The range of input colors that are mapped to visible colors (**0.0**Â â← Â **1.0**).

So without Exposure we will get a linear correction of all color values:

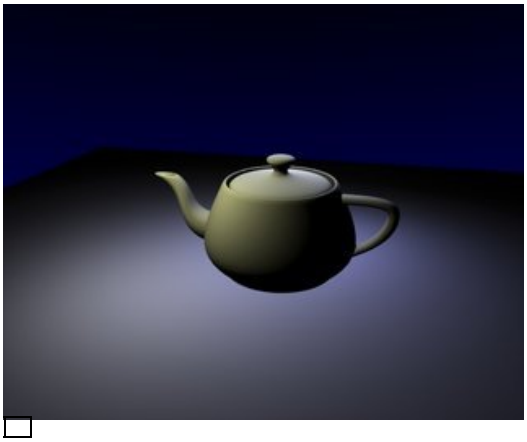
1. Range > **1.0**: the picture will become darker; with Range = **2.0**, a color value of **1.0** (the brightest by default) will be clipped to **0.5** (half bright) (*Range: **2.0***).
2. Range < **1.0**: the picture will become brighter; with Range = **0.5**, a color value of **0.5** (half bright by default) will be clipped to **1.0** (the brightest) (*Range: **0.5***).

Examples

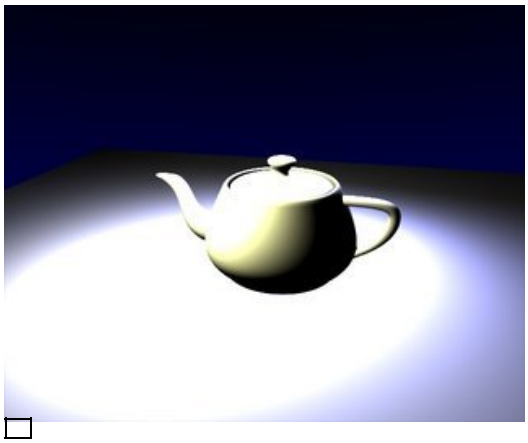
With a linear correction every color value will get changed, which is probably not what we want. Exposure brightens the darker pixels, so that the darker parts of the image wonâ← t be changed at all (*Range: **2.0**, Exp: **0.3***).



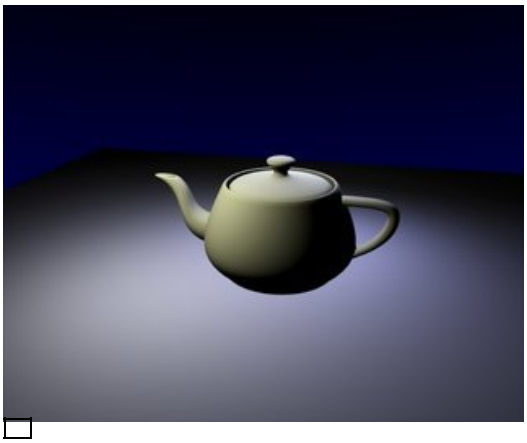
An overexposed teapot.



Range: **2.0**.



Range: **0.5**.



Range: **2.0**, Exposure: **0.3**.

Hints

Try and find the best Range value, so that overexposed parts are just not too bright. Now turn up the Exp value, until the overall brightness of the image is satisfying. This is especially useful with area lamps.

Introduction to Shading

The purpose of this section is to guide you in understanding the full array of Blender's shading options and how to use them. Shading is the process of applying color, textures, and finishes to your meshes in order to simulate a wide variety of appearances, including patterns, actual painting and detailing, faces of people and animals in a variety of settings.

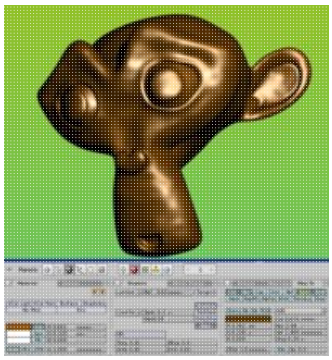
While Blender does offer painting abilities, those abilities are geared toward animation or coloring meshes. At a fine level of detail, such as skin UV textures or matte paintings, it does not compete and does not hope to compete with ease of use and functionality in specialized paint programs such as Gimp or Photoshop. Instead, it seamlessly uses their output graphic files for mesh and scene colorization.

Basics: Materials and Textures

In the next sections, we will discuss Procedural Materials and Textures. These types of materials and textures are applied as a procedure across the entire mesh. In Multiple Materials we will discuss how to apply a material to a set of faces. Since a mesh can have many sets of faces, a mesh can have many procedural materials, each with its own color and other material settings, as well as procedural textures.

In Textures, we will see how textures such as clouds or wood grain can have subtle or dramatic effects on the appearance of a mesh, and how they overlay or enhance the base material in some way (depending on how they Map Output) and mix with previous textures.

Advanced: Vertex painting and UV texturing



Chocolate Monkey by Roger

After the basics, we will take our shading to the next levels: vertex painting and then UV texturing using images. Please note that these levels of complexity and detail work together. Images can be combined with vertex colors to make the texture brighter or darker or to give it a color. Using a special texture, the Image Texture, as a UV Texture is a way of giving very detailed, precise coloring to a mesh in order to achieve photo-realistic imagery.

Consider our favorite model, Suzanne the Monkey. Below you will see, in order of increasing realism (and complexity), the results that can be achieved with each style of painting.

In the picture to the right, a single material, brown, was used to color Suzanne. A cloud texture was applied to the Color to blend in a darker brown, as well as making it look bumpy (Normal) and to vary the specularity. Use this paint technique for an object that is physically made from a single material, like chocolate, plastic, or metal. Use one of the grain textures for wood or marble. Or a milk chocolate monkey; just eat the ears first.



Suzanne by Roger using Multiple Materials
In the picture to the right, a few materials, brown for fur, pink for nose and lips, etc. was used to color Suzanne. A cloud texture was applied to color her skin, and a veronni texture to her pink nose and eyes. A Vertex Group, Scalp, was defined to generate static particles (hair).



Mohawk Suzanne by RamboBaby using Vertex Painting
In the picture to the right, Vertex Painting was used to color Suzanne's mesh. A green scalp was used to generate a mohawk using static particles. Her right eye has the center vertex duplicated, set to single user, and a halo material was used to put "a twinkle in her eye". Her left eye has the cornea rotated to so that the vertex colors twist. Both eyes were filled in on the back side and vertex painted with green colors and also have Z-Trans enabled to allow the back to show through. Lights are placed in front of her eyes to highlight them.

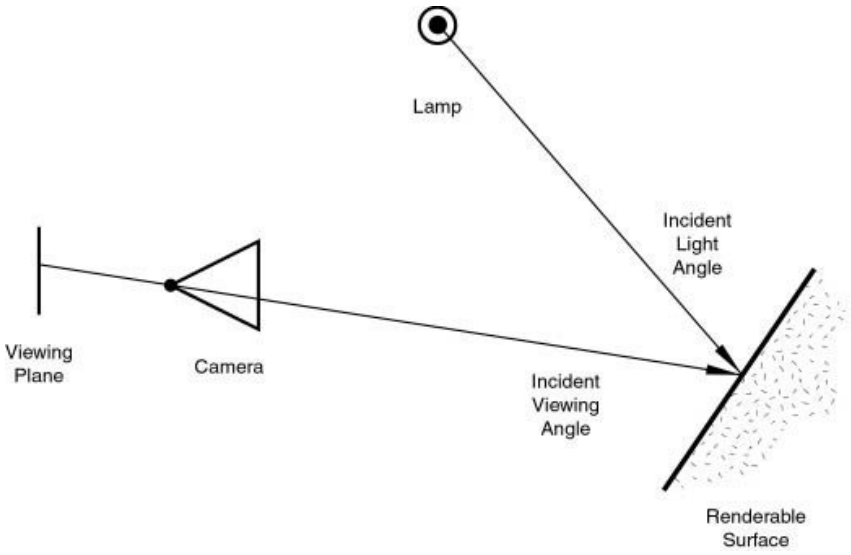
Since Vertex painting shades between faces, notice how the skin colors blend more smoothly than the previous example. Her skin has a more natural appearance with colors blended (and not being blocky or single-colored).

File:Manual-Painting-uvtexture.jpg

Suzanne by joe using UV Textures
In the picture to the right, a single image was used as a UV Texture to color Suzanne. A xxx texture was applied to xxx.

Introduction

Before you can understand how to design effectively with materials, you must understand how simulated light and surfaces interact in Blender's rendering engine and how material settings control those interactions. A deep understanding of the engine will help you to get the most from it. The rendered image you create with Blender is a projection of the scene onto an imaginary surface called the *viewing plane*. The viewing plane is analogous to the film in a traditional camera, or the rods and cones in the human eye, except that it receives simulated light, not real light. To render an image of a scene we must first determine what light from the scene is arriving at each point on the viewing plane. The best way to answer this question is to follow a straight line (the simulated light ray) backwards through that point on the viewing plane and the focal point (the location of the camera) until it hits a renderable surface in the scene, at which point we can determine what light would strike that point. The surface properties and incident light angle tell us how much of that light would be reflected back along the incident viewing angle (*Rendering engine basic principle.*).



Rendering engine basic principle.
Two basic types of phenomena take place at any point on a surface when a light ray strikes it: diffusion and specular reflection. Diffusion and specular reflection are distinguished from each other mainly by the relationship between the incident light angle and the reflected light angle. The shading (or coloring) of the object during render will then take into account the base color (as modified by the diffusion and specular reflection phenomenon) and the light intensity

Using the internal raytracer, other (more advanced) phenomena could occur. In raytraced reflections, the point of a surface stroke by a light ray will return the color of its surrounding environment, according to the rate of reflection of the material (mixing the base color and the surrounding environment's) and the viewing angle. On the other hand, in raytraced refractions, the point of a surface stroke by a light ray will return the color of its background environment, according to the rate of transparency (mixing the base color and the background environment's along with its optional filtering value) of the material and the optional index of refraction of the material, which will distort the viewing angle.

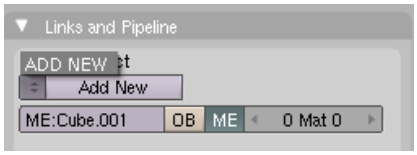
Of course, shading of the object hit by a light ray will be about mixing all these phenomena at once during the rendering. The appearance of the object, when rendered, depends on many inter-related settings:

- World (Ambient color, Radiosity, Ambient Occlusion)
- Lights
- Material settings (including ambient, emission, and every other setting on every panel in that context)
- Texture(s) and how they are mixed
- Material Nodes

- Camera
- viewing angle
- obstructions and transparent occlusions
- shadows from other opaque/transparent objects
- Render settings
- Object dimensions (SS settings are relevant to dimensions)
- Object shape (refractions, fresnel effects)

In this section we look at how to set up the various material parameters in Blender, and what you should expect as a result. We also give hints about practical material usage.

Creating a new Material



Add new material.
Every time a new Object is created it has no material linked to it. By pressing the F5 key or clicking on , you switch to the Shading context and the Material Buttons window appears. This window should be almost empty at this point.

Adding a new material is done with the menu button shown in *Add new material.*. The following options are available:

Add New
: Add a new material and link it to the active object or object data. Like other datablocks, Blender will automatically set its name to Material.001 and so on. It's a very good idea to give your materials clear names so you can keep track of them, especially when they're linked to multiple objects.

Select an existing material ()
: Choose an existing material from a list. If there are LOTS of materials, you will see a choice "DataSelect". Click that and one of your windows will change to a data browser window, listing all of the materials for you to choose from.

Sharing a Material from another object

Blender is built to allow you to reuse *anything*, including material settings, between many objects. Instead of creating duplicate materials, you can simply re-use an existing material. There are two ways to do this:

- With the mesh selected, click the up-down selector arrow located to the left of the Material name. The popup list shows you all the current materials. To use one, just click on it.
- In the 3D View, with CtrlL you can quickly link all selected objects to the material (and other aspects) of the active object. Very useful if you need to set a large number of objects to the same material; just select all of them, then the object that has the desired material, and CtrlL link them to that "parent".

Materials Options

Once the object has at least one material linked to it, many new panels (some tabbed) are readily displayed to allow you to precisely control the shading of the material. Using each of these panels is discussed in the next

section.

The preview panel attempts to show you what the shader will produce for different kinds of geometric basic shapes. Depending on your buttons window layout, some panels may be tabbed under others, or collapsed.

Material Preview

Mode: All Modes

Panel: Shading/Material Context → Preview

Hotkey: F5

Description

The Preview panel gives a quick visualisation of the active material and its properties, including its Shaders, Ramps, Mirror Transp properties and Textures. It provides many useful shapes that are very useful for designing new shaders: for some shaders (like those based Ramp colors or a Diffuse shader like Minnaert), one needs fairly complex or specific previewing shapes to decide if the shader-in-design achieves its goal.

Options

Flat XY plane
: Useful for previewing textures and materials of flat objects, like walls, papers and such.

Sphere
: Useful for previewing textures and materials of sphere-like objects, but also to design metals and other reflective/transparent materials, thanks to the checkered background.

Cube
: Useful for previewing textures and materials of cube-like objects, but also to design procedural textures. Features a checkered background.

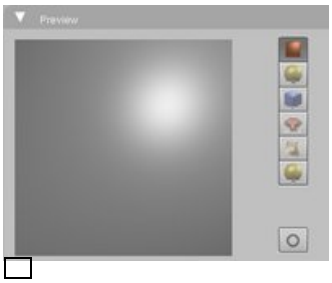
Monkey
: Useful for previewing textures and materials of organic or complex non-primitive shapes. Features a checkered background.

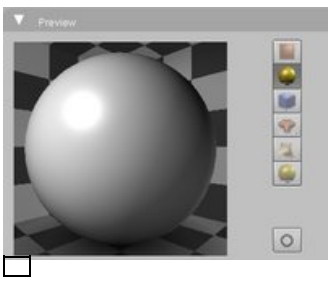
Hair strands
: Useful for previewing textures and materials of strand-like objects, like grass, fur, feathers and hair. Features a checkered background.

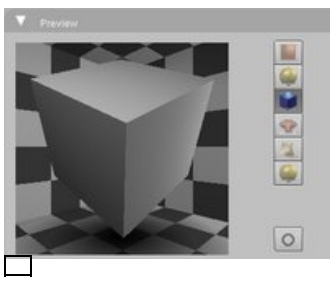
Large Sphere with Sky
: Useful for previewing textures and materials of sphere-like objects, but also to design metals and other reflective materials, thanks to the gradient Sky background.

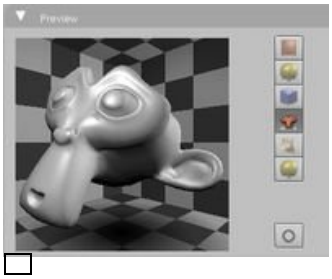
Preview uses OSA (oversampling)
: Whatever the preview option, it will make use of OSA (oversampling) in order to provide with a better quality. Disable this option if your computer is already slow or old.

Examples

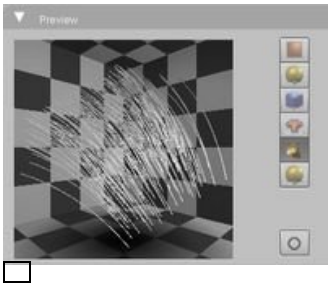

Plane preview.


Sphere preview.


Cube preview.



Monkey preview.



Hair Strands preview.



Sky Sphere preview.

Materials

Mode: All Modes

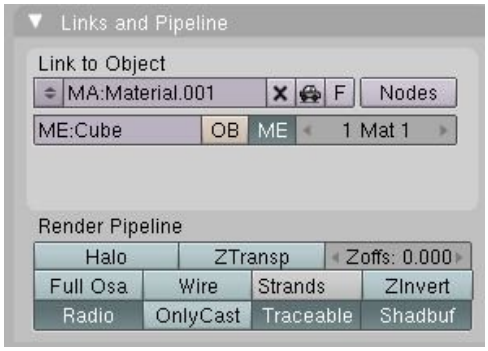
Panel: Shading/Material Context → Material

Hotkey: F5

Description

Materials can be linked to objects and obData in the materials panel, of the Shading/Material context. Here is where you can manage how materials are linked to objects, meshes, etc. and activate a material for editing in the rest of the panels.

Links and Pipeline Options



Links and Pipeline Tab
With a material linked or created, further options are available:

MA
(Material name, 

)
Shift-click into this field to name your material

Number of users (number field)
The number of objects or obdata that use the material. This material is linked between the various objects, and will update across all of them when edited. Clicking this number will make a 'single user copy', duplicating the material, with it linked only to the active object/obdata.

Auto ()
Automatically generates a (hopefully relevant) name for the material, based on the diffuse colour

F (Fake user)
Gives the material a 'fake user', to keep the material datablock saved in the .blend file, even if it has no real users

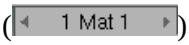
Nodes
Designates this material to be a material node noodle, and not from the Material/Ramps/Shaders settings.

Datablock links (

)
These two buttons determine whether the material is linked to the object or to (in this case) the mesh (or curve, nurbs, etc.). The ME button determines that this material will be linked to the mesh's datablock which is linked to the object's datablock. The OB button determines that the material will be linked to the object's data block directly.

This has consequences of course. For example, different objects may share the same mesh datablock. Since this datablock defines the shape of the object any change in edit mode will be reflected on all of those objects. Moreover, anything linked to that mesh datablock will be shared by every object that shares that mesh. So, if the material is linked to the mesh, every object will share it.

On the other hand, if the material is linked directly to the object datablock, the objects can have different materials and still share the same mesh. Short explanation: If connected to the object, you can have several instances of the same obData using different materials. If linked to mesh data, you can't.

Material indices (

)

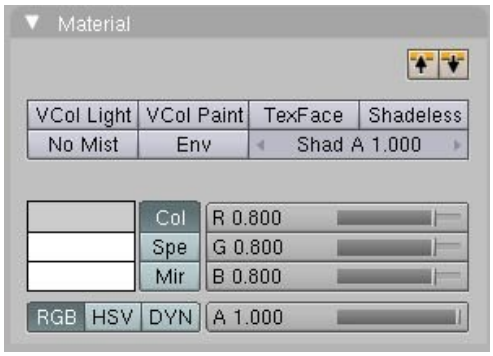
This shows how many materials you have for that obData and which one is currently active for editing. You can have multiple materials on an object and this can be done in the Editing Context (F9) in the Link and Materials Tab. See Doc:Manual/Materials/Multiple_Materials for more info.

Render Pipeline

These toggles tell Blender where this material fits into the Render Pipeline, and what aspects of the material are to be rendered.

- ◊ Halo - Each vertex is a halo of light
- ◊ ZTransp - Objects behind this one can shine through it
- ◊ Zoffs: A numeric input for offsetting the z-value; negative numbers puts this surface just in front of others assigned to the same mesh
- ◊ Full OSA - perform full over-sampling on this material to make it as smooth as possible where it overlays or intersects with other materials.
- ◊ Wire - render it in wireframe mode
- ◊ Strands-Colors static particles (hair, fur)
- ◊ ZInvert-renders faces inside-out. With OSA on, the render of one object against the backdrop of another object may occur *after* the render of the first object, and therefore the OSA calculation may pick up a halo or outline of the original backdrop. Use this option to disable this.
- ◊ Radio-material participates in Radiosity baking
- ◊ Only Cast - Material does not show up as a color, but just casts shadows onto other objects.
- ◊ Traceable-participates in ray tracing, if enabled during render
- ◊ Shadbuf-Can participate in shadow buffering, a faster alternative to raytracing

Material Options



Material Tab

Copy / Paste (
)

Copies the active material's settings / Pastes stored material settings into the active material

Modes

VCol Light: kinda like Halo, but each vertex's color is used as a light source

VCol Paint: enables Vertex Paint

TexFace: colors are taken from a UV Texture

Shadeless - color is not affected by light or shadows
No Mist - color is not obscured or faded by world mist
Env - renders it invisible
Shad A - this numeric control sets the Alpha value of shadows cast by this mesh, thus softening them

Colors

There are three colors to set:

◊ Col - Diffuse Color
◊ Spe - Specular Color
◊ Mir - Mirror Color.

Clicking a color swatches brings up the color picker applet, which allows you to pick a pre-defined standard color: to learn about it, please read the Color Picker reference.
You can also click on the eyedropper to sample a color, or play with the RGB sliders, or play with the HSV box and click on a color inside the gradient. The RGBA sliders set the active color. Click Col, Spe, or Mir to activate a color settings for the sliders to set.

Color Expression

RGB is for Red Green Blue sliders;
HSV changes the sliders to set Hue Saturation Value

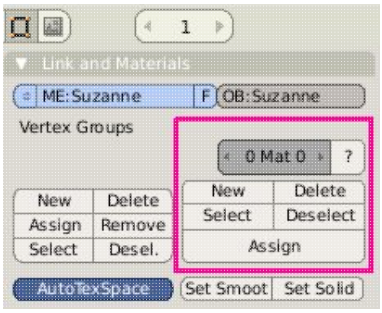
Dynamics

For the Game Engine, these set up friction for the surface

We normally think of an object as having a single material; a color, maybe with a little texture. This works for modeling many of the simpler objects in the real world. However, the real world is not always that simple, and some objects get very complicated. Blender allows you to have multiple materials for different pieces of a mesh. This section describes how to use the multiple material features of Blender.

To begin, make sure your 3D window is in Solid, Shaded, or Textured mode so that you can see the effects of your changes.

The Material Index



Default Material Index display

The list of materials that are assigned to a mesh is located by viewing the Editing buttons F9 in a Buttons Window on the Link and Materials tab. On that tab, the buttons used for defining and managing multiple materials are boxed in hot pink in the image to the right. These buttons show the active material name (currently blank), a Material Index display (the 0 Mat 0 ? scroller), and some action buttons labeled New, Delete, Select, Deselect, and Assign. The default panel shown to the right shows that the object Suzanne does not have any materials assigned to it, as evidenced by the 0 Mat 0 material index display.

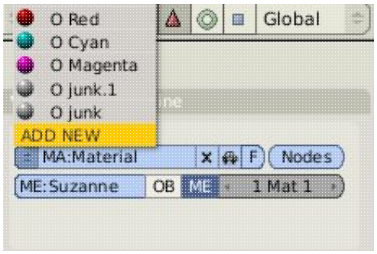


Material Index:

In the < 5 Mat 2 > display, the two small arrows allow you to change the active material - if the object uses more than one of them; the first number (**5** here) gives you the number of materials used by the current object, and the Mat 2 indicates that it's the material of index **2** which is active.

If a mesh does not have any materials (0 Mat 0 is showing), create a quick and simple new material for the object by clicking the New button. A new material will be created and assigned to the object. It's name, "Material" or "Material.00x", will be shown above the Index display. By default, it will be grey. The Index will now display 1 Mat 1

Meshes and Objects
You can re-use the same mesh within Blender many times to create objects. For example, you may create a mesh called "chair" and then create four duplicate objects and put them around a table. Each instance of the mesh, namely each object, may have its own set of materials assigned (one red chair, one blue chair, ...)



Shading Material panel Index display
The **other** place the material index for an object is shown is on the Links and Pipeline panel on the Shading F5 Material Buttons window. The image to the right shows that our Suzanne has one MA:terial called "Material" assigned to it, and that it is the first of one materials assigned to her (the 1 Mat 1 index display). The select/scroll button has been clicked and is showing the list of materials in the blend file, and the ADD NEW option is highlighted. You can use this option to create a new material for the object as previously discussed in preceding sections.

Some of the mesh faces can be assigned to different index. Some of the indexes can use the same material. Changing the material settings (e.g. color) for a material will change the color of all the faces in all the indexes that use that color. To make it even more flexible (and confusing), the same Material can be used by many different meshes in your scene, so changing a material may even change the color of other meshes. In other words, be careful when changing material colors so that you don't have any unwanted effects. If you are ever in doubt, use the Outliner window to see who uses a particular material.

To make different indexes use different colors, you have to Add New materials, one for each of the indexes, as discussed below. There are three steps, and the steps can be done in any order:

- Creating Material Index slots
- Assign Faces to a Material Index
- Associate a Material with a Material Index

Assigning Material Indexes to a Mesh

A mesh can have multiple materials assigned to it, and the same material can be re-used many times for many meshes. Blender keeps track of which materials are used with which meshes by keeping a list of materials that a mesh uses. Each material is assigned an index, or slot, on the particular mesh's list.

Adding a Material Index

By clicking the New button, you create a new entry or slot on the object's list. Each time you click the New button, Blender adds a new material to the active index.

Active Material
On the Shading Material buttons panels, you are working with a particular material. The material that you are working on is called the Active Material.



Material panel showing 4 indices; index 2 active
The image to the right shows our Editing F9 panel again, this time telling us that Suzanne now has 4 slots, and that we are working with the second, Cyan, as the active material. The color of the material is shown in the swatch.

To select a Material Index, use the left/right scroll arrows (shown in the image as little arrows to the left and right of the 4 Mat 2).

The main, or default material for the object is always in slot 1.

If you want to use an existing material for the new index, you must be in edit mode and use the pulldown menu that will appear in the panel. See this thread for details.

Deleting a Material Index

Take your object out of edit mode. Select the material index slot using the scroll buttons in the material Index display. To delete this index, click the Delete button. Clicking this button disassociates the material with the mesh. It does not delete the material definition from the .blend file, but simply crosses it off the list for that particular object. Any faces of the mesh that may have used that material index will be re-assigned the default material (index/slot 1).

Assigning Face(s) to a Material Index

A "face" is the part of the mesh that actually reflects color and shows the material. A mesh is made up of many faces, connected by edges with corners defined by vertices. All faces of the mesh are assigned to the first material index (x Mat 1) when the mesh was created.

With your object in Edit mode (â Tab), select the faces that you want to have a particular color (material) in the 3D window. Recall that there are vertex, edge, and face select buttons in the window header. You can select multiple faces using â § ShiftRMB or Border select. Then, in the Material Index panel in the Buttons window, scroll or ensure that the correct index is showing in the x Mat x material index scroller. Press the Assign button to assign those selected faces to the active material index.

Keeping Track
Unlike Vertex Groups, the Material Indexes cannot be named. So, you might want to open a Text Editor window and jot down what Indexes are used for what faces. For example, 1 for base, 2 for ears, 3 for eyes, 4 for nose.

To assign a face to a different index, simply select the face, scroll to the material you want it to show, and press Assign. The face was assigned to a different index, and that assignment is overridden to what you just specified. If you made a mistake, there's always Undo.

Changing a Material used by a Material Index

When you clicked the New button in the Material Index panel, the active material was assigned to that index. At that moment, the previous index *and* the new index both pointed to and used the *same* material. Since they both used the same material, nothing changed in your display. However, you probably want them to be different.

Use the Shading Material buttons to change the appearance of an active material that is referenced by material indexes. In the Buttons window, the Shading (F5 set), Material buttons, Links and Pipeline panel, use the material index Literal scroller to scroll to the index you want to be a different color. Then use the MA: material scroll/select button to Add New material. All of the faces assigned to that index will now show the new color in your 3D window.

To swap out a material on the list for another, go to the Shading Material buttons, Links and Pipeline panel. Use the MA: material scroll/select button to select a different material for that index (recall the panel also displays the active material in its own x Mat x index display). All of the faces assigned to that index will now have the new color.

Material Users
In the Material select scroll box that shows all the materials in your scene, a number to the right of its name shows the number of objects that share that material. A zero means that the material is unused and readily available without affecting any other objects or faces.

Checking your work using Select, Deselect, Query (?)

To see which faces have which index, in Edit mode, deselect all faces by pressing the A key once or twice. Then click the Select button. In the 3D window, the faces that have the active material will be selected, just as if you selected them manually using RMB .

If you see that a face has the wrong material assigned to it, select it using the RMB , scroll to the correct material index, and press the Assign button.

To see if any faces do NOT have a material assigned to them, use the Deselect button as follows:

- Select all the faces of the object (A once or twice)
- Scroll to the first material index, and click Deselect. All the faces that are assigned that index will be de-selected.
- Repeat the above step for each index.
- Any faces left are unassigned a material index.

To see the material that a face is assigned to, select the face and click the ? question mark query button next to the x Mat x Material Index display. The display will change to show that face's material index.

Some Examples



A soda can has base aluminum, but an image wrapper. The pull tab may be a colored aluminum (three indexes). A chair may have a frame and a cushion (two materials). A old-fashioned pencil has an eraser, metal holder, painted wood stylus, shaved part and graphite (five materials). Sunglasses have a frame and lenses (two materials). The computer I am writing this on has about ten. Our favorite flag (*right*) is one object with three materials (orange, blue and white) The white material index is the default (index 1) and is assigned to both the stars and the background.

Advanced Multiple Materials

A way of mixing materials on a mesh without using geometry for separation is shown in the Mixing Materials with Nodes tutorial.

Diffuse Shaders

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5

Description

A diffuse shader determines, simply speaking, the general colour of a material when light is shined on it. Most shaders that are designed to mimic reality give a smooth falloff from bright to dark from the point of the strongest illumination to the shadowed areas, but Blender also has other shaders for various special effects.

Options

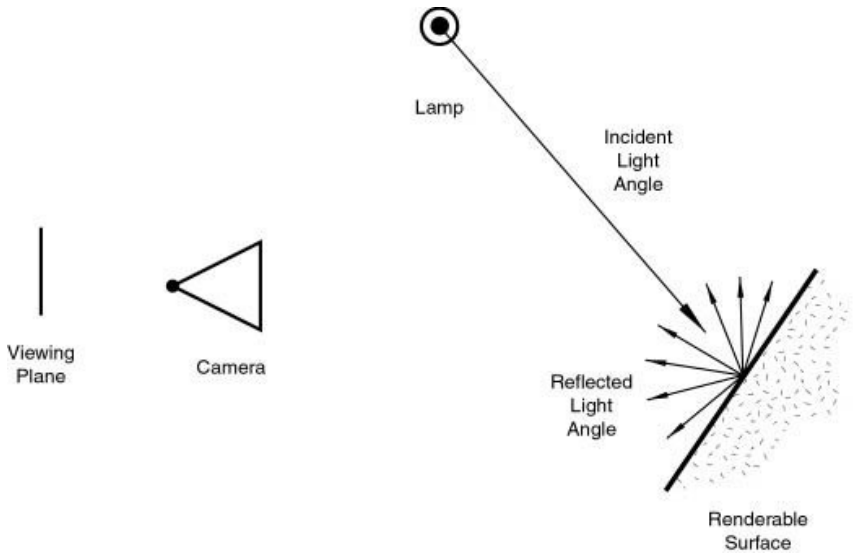
All diffuse shaders have the following options:

Color
: The base *Diffuse color* of the material

Ref
: The shader's brightness, or more accurately, the amount of incident light energy that is actually diffusely reflected towards the camera.

Technical Details

Light striking a surface and then re-irradiated via a Diffusion phenomenon will be scattered, i.e., re-irradiated in all directions isotropically. This means that the camera will see the same amount of light from that surface point no matter what the *incident viewing angle* is. It is this quality that makes diffuse light *viewpoint independent*. Of course the amount of light that strikes the surface depends on the incident light angle. If most of the light striking a surface is reflected diffusely, the surface will have a matte appearance (*Light re-irradiated in the diffusion phenomenon.*).



Light re-irradiated in the diffusion phenomenon.

Hints

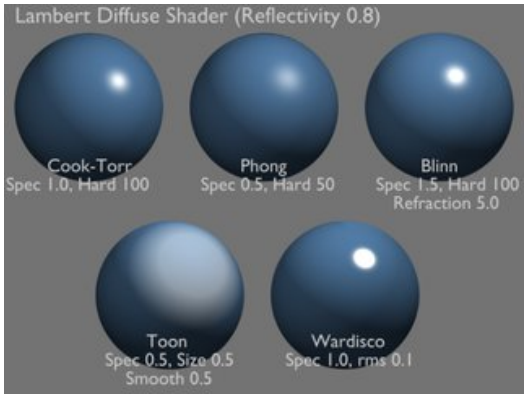
Some shaders' names may sound odd - they are traditionally named after the people who invented them.

Lambert

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5

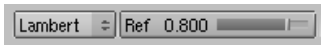


Lambert Shader

Description

This is Blender's default diffuse shader, and is good general all-around work horse...neigh!

Options



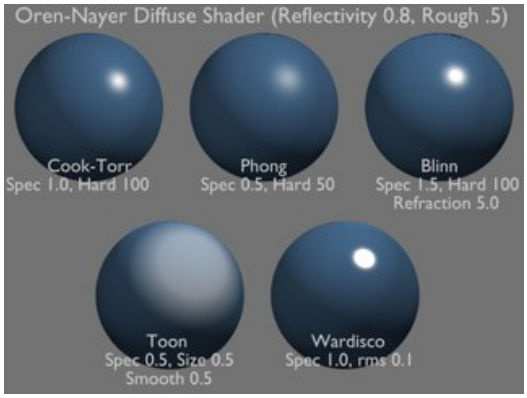
The Lambert diffuse shader settings.
This shader has only the default option, determining how much of available light is reflected. Default is 0.8, to allow some other objects to be brighter.

Oren-Nayar

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5

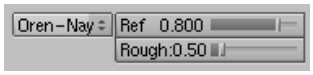


Oren-Nayer Shader

Description

Oren-Nayar takes a somewhat more 'physical' approach to the diffusion phenomena as it takes into account the amount of microscopic roughness of the surface.

Options



The Oren-Nayar diffuse shader settings.

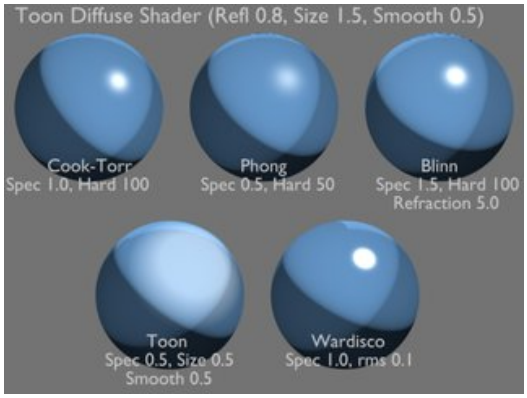
Rough
: The roughness of the surface, and hence, the amount of diffuse scattering

Toon

Mode: All Modes

Panel: Shading/Material Context → Shaders

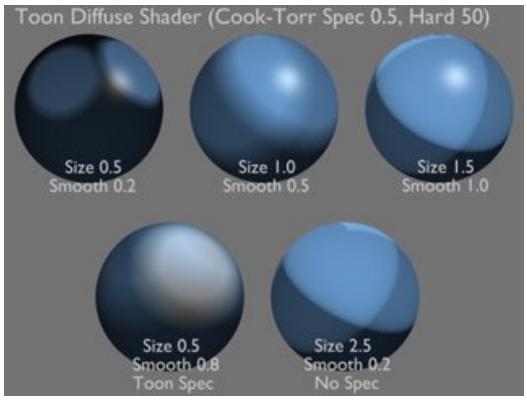
Hotkey: F5



Toon Shader, Different Spec

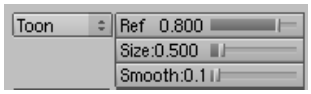
Description

The Toon shader is a very 'un-physical' shader in that it is not meant to fake reality but to produce cartoon cel styled rendering, with clear boundaries between light and shadow and uniformly lit/shadowed regions.



Toon Shader Variations

Options



The Toon diffuse shader settings.

Size
: The size of the lit area

Smooth

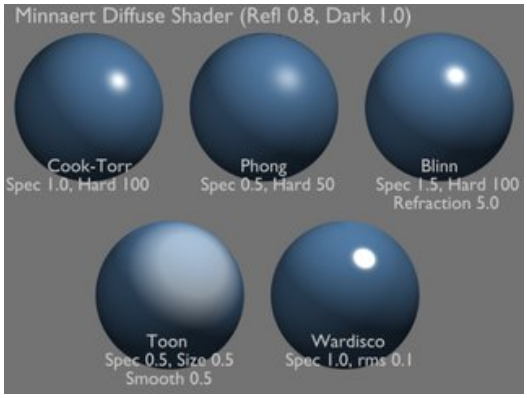
The softness of the boundary between lit and shadowed areas

Minnaert

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5

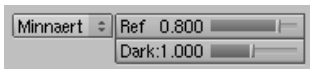


Minnaert Shader

Description

Minnaert works by darkening parts of the standard Lambertian shader, so if *Dark* is 1 you get exactly the Lambertian result. Higher darkness values will darken the center of an object (where it points towards the viewer). Lower darkness values will lighten the edges of the object, making it look somewhat velvet.

Options



The Minnaert diffuse shader settings.

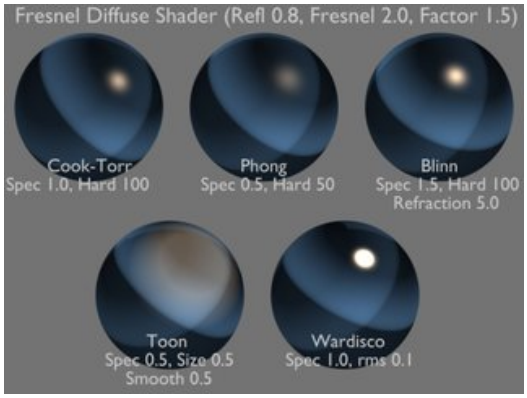
Dark
: The darkness of the 'lit' areas (higher) or the darkness of the edges pointing away from the light source (lower).

Fresnel

Mode: All Modes

Panel: Shading/Material Context → Shaders

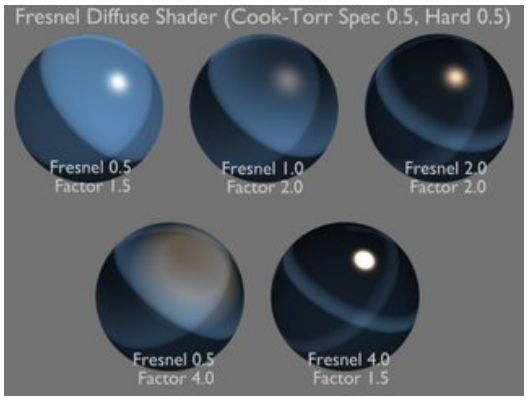
Hotkey: F5



Fresnel Shader, Different Spec

Description

With a Fresnel Shader the amount of diffuse reflected light depends on the incidence angle, i.e. from the direction of the light source. Areas pointing directly towards the light source appear darker, areas perpendicular to the incoming light become brighter.



Fresnel Shader, Same Spec

Options

Ref is the Reflectivity; amount of color reflected for each unit of light received. Fresnel is the power of the Fresnel effect, and Fac is the amount of the effect to blend in.

Other Options

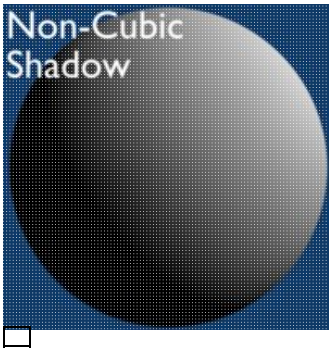
If you look farther down the side of the shaders tab, you will see several other buttons.

- Cubic:

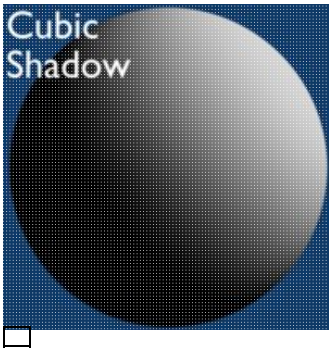
 Uses cubic interpolation of diffuse values for smoother transitions (between light and dark).

- Bias:

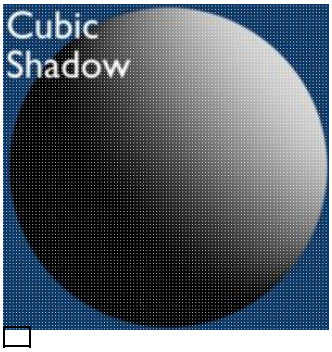
 Eliminates ray-traced shadow errors with the Phong specular shader.



Without Cubic enabled.



With Cubic enabled.



Animation switching between Non-Cubic and Cubic shadowing. You will need a modern, standards compliant, browser to see the animation. Click to View Animation.

Specular Shaders

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5

Description

Specular shaders create the bright highlights that one would see on a glossy surface, mimicking the reflection of light sources. Unlike diffuse shading, specular reflection is *viewpoint dependent*. According to Snell's Law, light striking a specular surface will be reflected at an angle which mirrors the incident light angle (with regard to the surface's normal), which makes the viewing angle very important.

Note
It is important to stress that the *specular reflection* phenomenon discussed here is not the reflection we would see in a mirror, but rather the light highlights we would see on a glossy surface. To obtain true mirror-like reflections you would need to use the internal raytracer. Please refer to section RENDERING of this manual.

Options

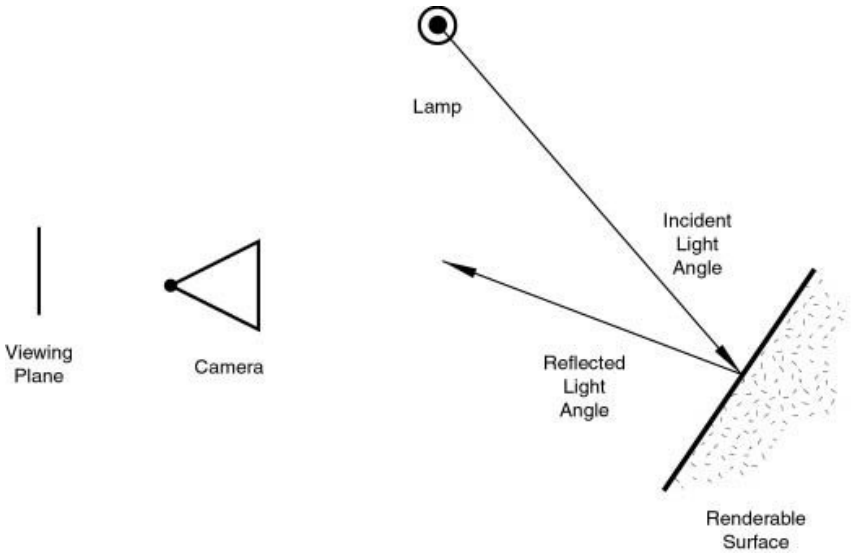
Each specular shader shares two common options:

Specular colour
: The colour of the specular highlight

Spec
: The intensity, or brightness of the specular highlight. This has a range of [0-2], which effectively allows more energy to be shed as specular reflection as there is incident energy.

As a result, a material has at least two different colours, a diffuse, and a specular one. The specular color is normally set to pure white, but it can be set to different values for various effects.

Technical Details



Specular Reflection.
In reality, Diffusion and Specular reflection are generated by exactly the same process of light scattering. Diffusion is dominant from a surface which has so much small-scale roughness in the surface, with respect to wavelength, that light is reflected in many different directions from each tiny bit of the surface, with tiny changes in surface angle.

Specular reflection, on the other hand, dominates on a surface which is smooth, with respect to wavelength. This implies that the scattered rays from each point of the surface are directed almost in the same direction, rather than being diffusely scattered. It's just a matter of the scale of the detail. If the surface roughness is much smaller than the wavelength of the incident light it appears flat and acts as a mirror.

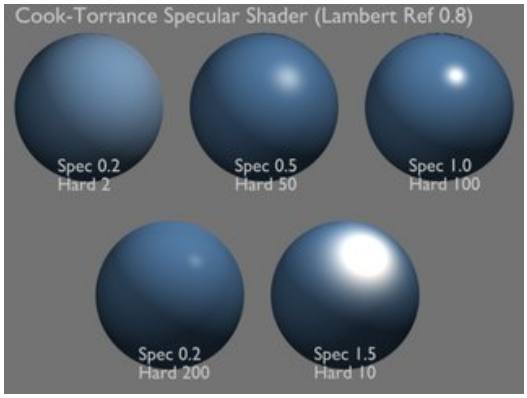
If it is difficult for you to understand the relation between roughness' scale and light's wavelength, try to imagine a ball (say, of centimetre scale): if you throw it against a wall of raw stones (with a scale of roughness of a decimetre), it will bounce in a different direction each time, and you will likely quickly lose it! On the other hand, if you throw it against a concrete wall (with a roughness of, say, a millimetre scale), you can quite easily anticipate its bounce, which follow (more or less!) the same law as the light reflectionâ←¦

CookTorr

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5



Cook-Torrance Shader

Description

CookTorr (Cook-Torrance) is a basic specular shader that is most useful for creating shiny plastic surfaces. It is a slightly optimised version of Phong.

Options

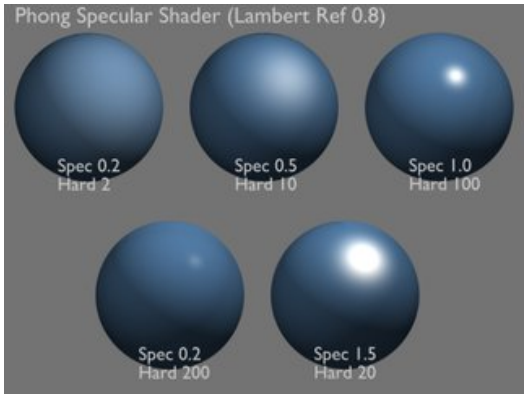
Hard
: The hardness, or size of the specular highlight

Phong

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5



Phong Shader

Description

Phong is a basic shader that's very similar to CookTorr, but is better for skin and organic surfaces

Options

Hard
: The hardness, or size of the specular highlight

Planet Atmosphere

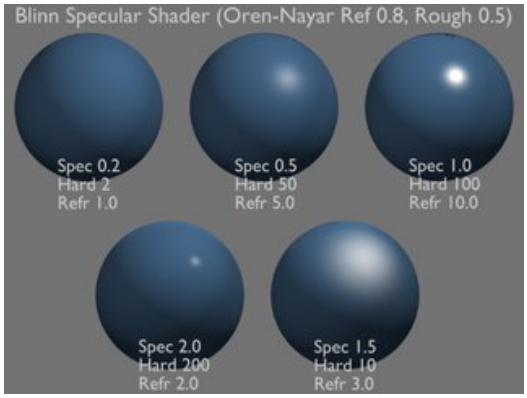
Because of its fuzzyness, this shader is good for atmosphere around a planet. Add a sphere around the planet, slightly larger than the planet. For its material, use a phong specular shader. Set it very low alpha (.05), zero diffuse, low hardness (5) but high specularity (2).

Blinn

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5



Blinn Shader

Description

Blinn is a more 'physical' specular shader, often used with the Oren-Nayar diffuse shader. It can be more controllable because it adds a fourth option, an *index of refraction*, to the aforementioned three.

Options

Hard
: The hardness, or size of the specular highlight. The Blinn shader is capable of much tighter specular highlights than Phong or CookTorr.

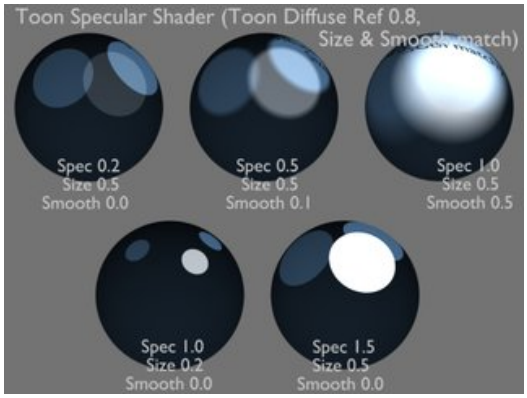
Refr
: The 'Index of Refraction'. This parameter is not actually used to compute refraction of light rays through the material (a ray-tracer is needed for that), but to correctly compute specular reflection intensity and extension via Snell's Law.

Toon

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5



Toon Specular Shader

Description

The Toon specular shader matches the Toon diffuse shader. It is designed to produce the sharp, uniform highlights of cartoon cels.

Options

Size
: The size of the specular highlight

Smooth
: The softness of the highlight's edge

Hints

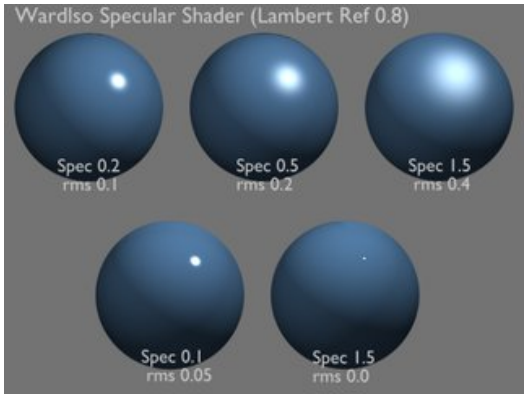
Toon shader can be also be accomplished in a more controllable way using ColorRamps.

WardIso

Mode: All Modes

Panel: Shading/Material Context → Shaders

Hotkey: F5



WardIso Shader

Description

WardIso is a flexible specular shader that can be useful for metal or plastic.

Options

rms
: rms controls, in effect, the size of the specular highlight, though using a different method to that of the other specular shaders. It is capable of extremely sharp highlights.

The following setting depends on how you set the Ambient Light.

Material Ambient Effect



The Ambient slider.
Each objectâ← s material setting (diffuse shader) has an Ambient slider that lets you choose how much ambient light that object receives. Generally, about 0.1 to half is good, depending on the color of the ambient light. A setting of 0.0 means that the object does not receive any ambient light; it is only lit by light that actually gets to it. A setting of 1.0 means that it is lit by all of the worldâ← s ambient light. Increasing the Ambient setting has the effect of flattening the shading, since the ambient light washes out the diffuse shading.

You should set this slider depending on the amount of ambient light you think the object will receive. Something deep in the cave will not get any ambient light, whereas something close to the entrance will get more. Note that you can animate this effect, to change it as the object comes out of the shadows and into the light.

Consequently, this slider also controls how much the material is affected by Ambient Occlusion. For instance, if you do not wish for the material to receive any Ambient Occlusion (lightening, darkening or both), move the slider all the way to the left-most (zero) position.

Ramps

Mode: All Modes

Panel: Context Shading â sub-context Material â Ramps

Hotkey: F5

In many real life situations - like skin or metals - the colour of diffuse and specular reflections can differ slightly, based on the amount of energy a surface receives or on the light angle of incidence. The new Ramp Shader options in Blender now allow you to set a range of colours for a Material, and define how the range will vary over a surface, and how it blends with the 'actual colour' (typically from material or as output of texture).

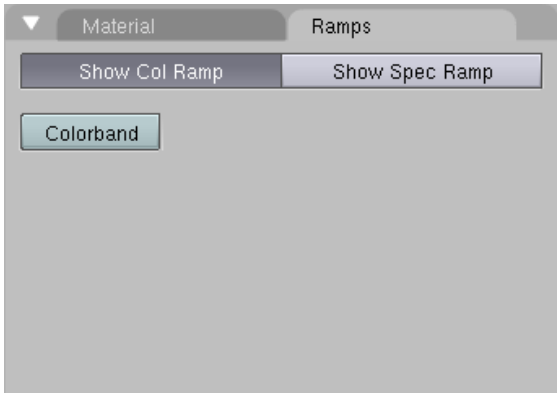
Description

Ramps allow you to precisely control the colour gradient across a material, rather than just a simple blend from a brightened colour to a darkened colour, from the most strongly lit area to the darkest lit area. As well as several options for controlling the gradient from lit to shadowed, ramps also provide 'normal' input, to define a gradient from surfaces facing the camera to surfaces facing away from the camera. This is often used for materials like some types of metallic car paint, that change colour based on viewing angle.

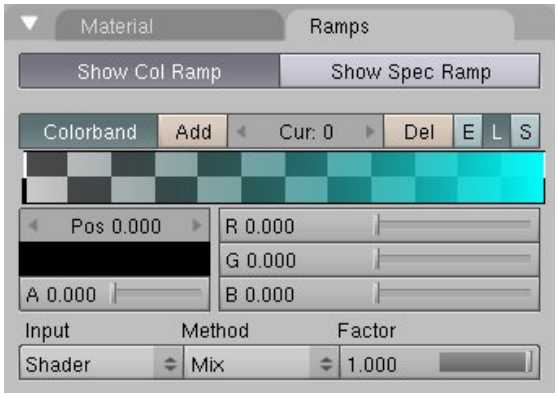
Since texture calculations in Blender happen before shading, the Ramp Shader can completely replace texture or material color. But by use of the mixing options and Alpha values it is possible to create an additional layer of shading in Blender materials.

Options

The Ramps panel is located in the Material context (F5). Here you can use the top two buttons to show either settings by pressing Show Col Ramp for diffuse or Show Spec Ramp for specular ramps (*Ramps Panel.*).



Ramps Panel.
Pressing the button Colorband enables Ramp Shaders. By default it opens with two colours, with the first having **Alpha = 0**, and no colour. And the second having **Alpha = 1** and a cyan colour (*Ramps Panel Colorband.*).



Ramps Panel Colorband.
See the settings description of the colorband below.

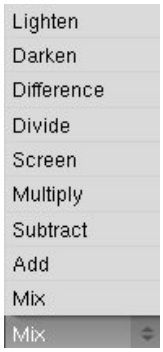
The two pop-up buttons and the value slider in the bottom of the panel defines how the Ramp Shaders work:



Input popup menu.

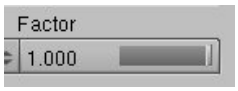
Input
: The input menu contains the following options for defining the gradient:
Shader
: The value as delivered by the material's shader (like Lambert or Phong) defines the colour. Here the amount of light doesn't matter for colour, only the direction of the light.
Energy
: As Shader, but now also lamp energy, colour and distance, is taken into account. This makes the material change colour when more light shines on it.
Normal
: The surface normal, relative to camera, is used for the Ramp Shader. This is possible with a texture as well, but added for convenience.
Result

All three previous options work per lamp, this option only does it in the end of all shading calculation. This allows full control over the entire shading, including 'Toon' style results. Using alpha values here is most useful for tweaking a finishing touch to a material.



Method popup menu.

Method
A list of the various blending modes available for blending the ramp shader with the colour from Input.



Factor slider.

Factor
The Factor slider denotes the overall factor of the Ramp Shader effect: **0.0** means no effect and **1.0** means full effect.

Colorbands

Mode: All Modes

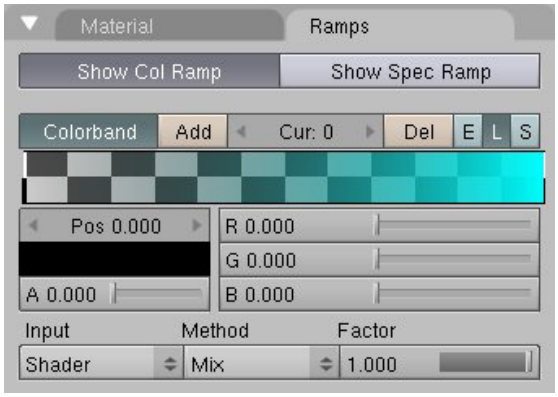
Panel: Context Shading â sub-context Material â Ramps

Hotkey: F5

Description

A colorband can contain a gradient through a sequence of many colours (with alpha), each colour acting across a certain position in the spectrum. Colorbands are used in both materials and textures, as well as other places where a range of colours can be computed and displayed.

Options



Ramps Panel Colorband.

Add
: Add a new mark to the centre of the colorband with default colours (neutral grey). New marks can also be added with CtrlLMB click on the colorband itself, which will add the mark at the position of the click, with the same colour as already existing underneath the mouse pointer.

Del
: Remove the current mark from the colorband.

Cur
: The number of the current selected mark on the colorband that's being edited. The current selected mark is shown on the colorband with double width. To select a colour mark you can either press LMB on the desired one, or step the current colour number up and down with the right and left arrows in this control. You can also â § ShiftLMB in the field and enter the required colour number manually.

Pos
: The location of the current mark along the colorband, from **0.0** (left) to **1.0** (right). The position of the marks can also be changed by clicking and dragging them with LMB.

Reordering Colours
If you reorder the positions of the colours, they will be renumbered so that they always start with **0** from the left and increment to the right.

R/G/B
: The colour of the current mark. You can LMB on the colour swatch under the Pos field to choose a colour using the Color Picker.

Alpha
: The alpha value (opacity) of the current mark. An Alpha value of **0.0** means that the colour is totally transparent and will not be seen in the final colorband. A value of **1.0** sets the colour opaque. If you defined colours with different Alpha values, they will be interpolated between each other to get a smooth transition between different transparency settings. You can preview the Alpha settings on the colorband with the checker board pattern behind the colorband. If pattern is visible then the transparency is less than **1.0**.

The colours can interpolate from one to the other in the following ways:

E
: Ease in by a quadratic equation.

C
: Cardinal.

L
: Linear (Default). A smooth, consistent transition between colours.

S
: B-Spline.

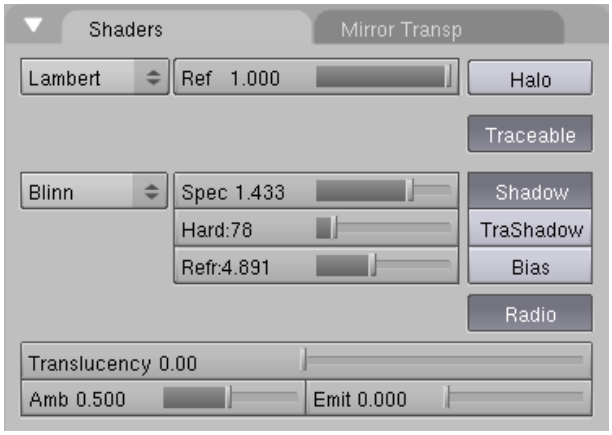
Examples



In this first example to the right, we want to texture a snake, specifically the deadly coral snake. We want to make a repeating set of four colours: black, yellow, red, yellow (and then back to black again). We also want to make the rings sharp in definition and transition. This example uses **8** colorband settings: **0** and **7** are black; **1** and **2** are yellow, **3** and **4** are red, and **5** & **6** are yellow. Position **0** and **1** close together, **2** and **3**, etc. Use a little noise and turbulence; together with the scales texture you should get really close!

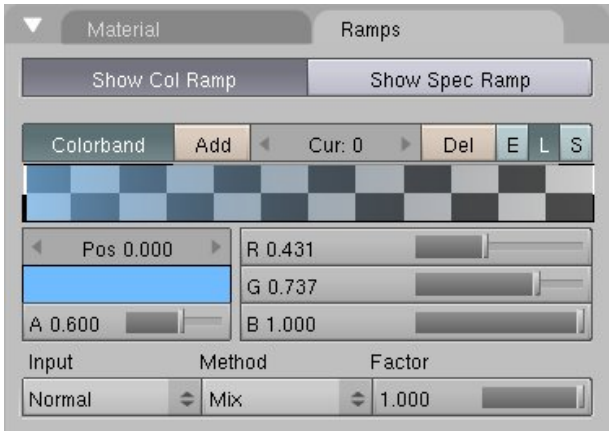
Let us make another simple example using Ramp Shaders.

- Remove default cube object from scene and add a Monkey mesh (â § ShiftA â Add â Mesh â Monkey). Press Subsurf (NOTE: In 2.42 and higher, Subsurf is now in the Modifier panels), and set display and render Subsurf Subdivision level to **2**. Press Set Smooth to get a nice smooth Monkey. All this in Edit context (F9).



Shader settings.

- Now press â Tab to exit Edit mode. Press F5 to enter Material context. In the Material panel press Add New to add a new material. Change the parameters in the Shaders tab as (*Shader settings*).



Ramp Shader settings.

- Press Ramps tab to open the Ramp Shader panel. Press Colorband to activate the Ramp Shader effect. Now try to set the parameters as (*Ramp Shader settings*). Remember to set the Input to Normal. The second colour to the right is set to Alpha = **0.0** and the colour to pure black.



Specular Ramp Shader color 0.

- In the Ramps tab press Show Spec Ramp and try to set the parameters as (*Specular Ramp Shader color 0*) and (*Specular Ramp Shader color 1*).



Specular Ramp Shader color 1.

Here is the rendered result of the settings we just entered above. In (*No Ramp Shader*) there is no Ramp Shader active. In (*Color Ramp*) the Color Ramp is activated and finally in (*Both Color and Specular Ramp*) both Color Ramp and Specular Ramp are activated! Remember here we have just demonstrated one effect of the Ramp Shader. There is much more to explore, try changing the Input and Method parameters, to see totally different results from the ones we have just seen in this example.



No Ramp Shader.



Colour Ramp.



Both Colour and Specular Ramp.

Raytraced Mirror Reflections

Mode: All Modes

Panel: Shading/Material Context → Mirror Transp

Hotkey: F5

Description

Raytracing can be used to make a material reflect its surroundings, like a mirror. The principle of raytraced reflections is very simple: a ray is fired from the camera and travels through the scene until it encounters an object. If the first object hit by the ray is not reflective, then the ray takes the color of the object. If the object is reflective, then the ray bounces from its current location and travels up to another object, and so on, until a non-reflective object is finally met and gives the whole chain of rays its color.

Eventually, the first reflective object inherits the colors of its environment, proportionally to its RayMir value. Obviously, if there are only reflective objects in the scene, then the render could last forever. This is why a mechanism for limiting the travel of a single ray has been set through the Depth value: this parameter sets the maximum number of bounces allowed for a single ray.

Note
You need to enable raytracing in your scene settings if you want to use raytraced reflections. This is done in the Scene/Render context → Render Panel. Raytracing is enabled by default in Blender 2.37 and higher.

The *Mir* setting on the Material panel is the color of the light reflected back. Usually, for normal mirrors, use White. However, some mirrors color the reflection, so you can change the color by clicking on the swatch on that panel. The Mirror Transp panel determines how and whether the mirror actually is active and reflects things. If you set RayMir to something >0, then the mirror is active. The reflection will be the tinted the color you set on the Material panel.

Options



The Mirror Transp Panel.

Ray Mirror

Enables and disables raytraced reflections

RayMir

Sets the amount of reflectiveness of the object. Use a value of 1.00 if you need a perfect mirror, or set RayMir to 0.00 if you don't want any reflection.

Fresnel

Sets the power of the Fresnel effect. The Fresnel effect controls how reflective the Material is, depending on the angle between the surface normal and the viewing direction. Typically, the larger the angle, the more reflective a Material becomes (this generally occurs on the outline of the object).

Fac

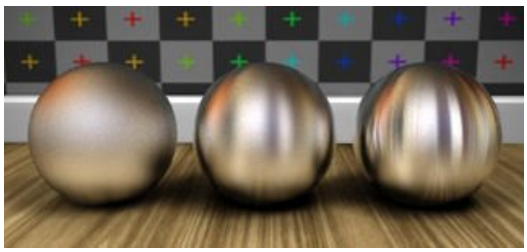
A controlling 'factor' to adjust how the blending (between reflective and non-reflective areas) happens.



Suzanne in the Fun House

Gloss

In paint, a high-gloss finish is very smooth and shiny. A flat, or low gloss disperses the light and gives a very blurry reflection. Also, uneven or waxed-but-grainy surfaces (such as car paint) are not perfect and therefore slightly need a Gloss < 1.0. In the example to the right, the left mirror has a Gloss of 0.98, the middle is Gloss = 1.0, and the right one has Gloss of 0.90. Use this setting to make a realistic reflection, all the way up to a completely foggy mirror. You can also use this value to mimic depth of field in mirrors.



Anisotropic tangent reflecting spheres with anisotropy of 0.0, 0.75, and 1.0

Anistropy

The amount a reflection is stretched in the tangent direction. If the tangent shading option is on, Blender automatically renders blurry reflections as anisotropic reflections.

When Tangent is switched on, the *Aniso* slider controls the strength of this anisotropic reflection, with a range of 1.0 (default) being fully anisotropic and 0.0 being fully circular, as is when tangent shading on the material is switched off. Anisotropic raytraced reflection uses the same tangent vectors as for

tangent shading, so you can modify the angle and layout the same way, with the auto-generated tangents, or based on the mesh's UV co-ordinates.

Samples
: The number of samples to average to arrive at the pixel's final color. More samples will give a smoother result, but the greater the number of samples, the slower the render speed.

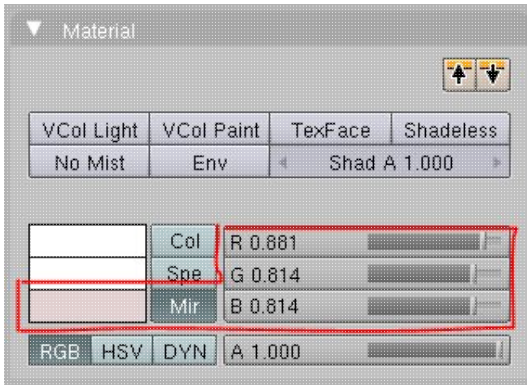
Threshold
: The threshold for adaptive sampling. Sampling is skipped when no more samples are deemed necessary, by checking the statistical variance of the samples so far for that pixel against the threshold. Raising the threshold will make the adaptive sampler skip more often, however the reflections could become noisier.

Depth
: Sets the maximum number of bounces for a single ray to be reflected. The default Depth of 2 is typically a good value. If your scene contains many reflective objects and/or if the camera zooms in on such a reflective object, you will need to increase this value if you want to see surrounding reflections in the reflection of the reflected object (!). In this case, a Depth of 4 or 5 is typically a good value.

Max Distance
: The number of blender units away from camera (Z-Depth) beyond which to stop calculating the actual reflection and to simply use the fade-out method. In the example, there is a mirror behind the camera, 10 blender units from the middle mirror. The top reflection of Suzy is actually a reflection of the reflection off the back mirror. The front mirror has Max Distance set to 20, so, as you can see, it is starting to fade to sky color.

Ray-End Fade Out
: For objects that recede away from the camera, further than Max Distance set above, you can have the mirror effect fade out, which reduces compute time. A large reflecting pond, at the far side, just fades out to be the sky color. You have two choices:
 - ◊ Fade to Sky color - Uses the sky color in the world settings
 - ◊ Fade to Material color - uses the material color

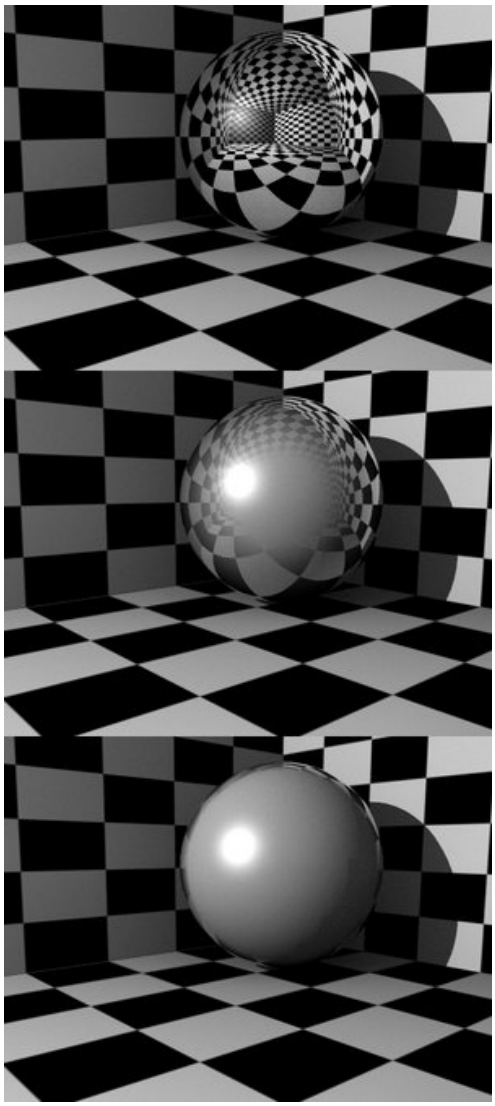
Colored reflections



Mirror Color
By default, an almost perfectly reflective Material like Chrome, or a Mirror object, will reflect the exact colors of its surrounding. But some other equally reflective Materials tint the reflections with their own color. This is the case for well polished copper and gold, for example. In order to replicate this within Blender, you have to set the Mirror Color accordingly. In the example above, the middle mirror has the mirror color set as shown to the right.

Examples

Fresnel



Demonstration of Fresnel effect with values equal to (from top to bottom) 0.0, 2.5 and 5.0

Let's undertake a small experiment in order to understand what Fresnel is really about. After a rainy day, go out and stand over a puddle of water. You can see the ground through the puddle. If you kneel just in front of the puddle, your face close to the ground, and look again at a distant point on the puddle of water, the liquid surface part which is closer to you lets you see the ground, but if you move your gaze towards the other end of the puddle, then the ground is gradually masked until all you see is the reflection of the sky. This is the Fresnel effect: having a surface sharing reflective and non-reflective properties according to the viewing angle and the surface normal.

In *Demonstration of Fresnel effect with values equal to (from top to bottom) 0.0, 2.5 and 5.0*, this behavior is demonstrated for a perfectly reflective Material (RayMir 1.0).

Fresnel 0.0 stands for a perfect mirror Material, while Fresnel 5.0 could stand for a glossy Material (varnished wood, for example?). It's barely noticeable but in the lower picture, the Material is perfectly reflective.

Reference

For more information, including other examples and in-depth discussion, please visit:

- Blender improved ray tracing
- QMC Sampling
- Mental Ray info

(note: all three links are broken)

Hints

In order to get a physically accurate Fresnel effect with the current algorithm, you have to set Fresnel to 5.0 and Fac to 1.25. Nevertheless, you can play with these values for the sake of artistic freedom, if you feel the need to.

Environment Maps (EnvMaps) can also be used to simulate reflective materials. Environment maps are more complicated to set up, have many limitations and are much less accurate, particularly on non-planar surfaces. However, Environment maps can be a lot faster to render and support extra features like filtering the reflection map to fake blurred reflections.

Raytraced Transparency

Mode: All Modes

Panel: Shading/Material Context → Mirror Transp

Hotkey: F5

Description

Raytracing is also used for simulating the refraction of light rays through a transparent material, like a lens. A ray is sent from the camera and travels through the scene until it encounters an object. If the first object hit by the ray is non-transparent, then the ray takes the color of the object. If the object is transparent, then the ray continues its travel through it to the next object, and so on, until a non-transparent object is finally encountered which gives the whole chain of rays its color. Eventually, the first transparent object inherits the colors of its background, proportionally to its Alpha value (and the Alpha value of each transparent Material hit in-between). But while the ray travels through the transparent object, it can be deflected from its course according to the Index of Refraction (IOR) of the material. When you actually look through a plain sphere of glass, you will notice that the background is upside-down and distorted: this is all because of the Index of Refraction of glass.

Enable Raytracing
You need to enable raytracing in your scene settings if you want to use raytraced transparency and refraction. This is done in the Scene/Render context → Render Panel. Raytracing is enabled by default in Blender 2.37 and higher.

Alpha value
You need to set your Alpha value (A in the Material Panel) at something else than 1.000. However, the alpha value will still be reported in the render as 1.0

Options



The Mirror Transp Panel.

RayTransp
: Enables and disables raytraced transparency

Filt
: Amount of filtering for transparent ray trace. The higher this value, the more the base color of the material will show. The material will still be transparent but it will start to take on the color of the material.

IOR
: Sets how much a ray travelling through the Material will be refracted, hence producing a distorted image of its background. See Examples.

Depth
: Sets the maximum number of transparent surfaces a single ray can travel through. There is no typical value. Transparent objects outside the Depth range will be rendered pitch black if viewed through the transparent object that the Depth is set for. In other words: if you notice black areas on the surface of a transparent object, the solution is probably to increase its Depth value (this is a common issue with raytracing transparent objects). You may also need to turn on transparent shadows on the other object.

Limit
: Materials thicker than this are not transparent. This is used to control the threshold after which the filter color starts to come into play.

Falloff
: How fast light is absorbed as is passes through the material. Gives 'depth' and 'thickness' to glass.

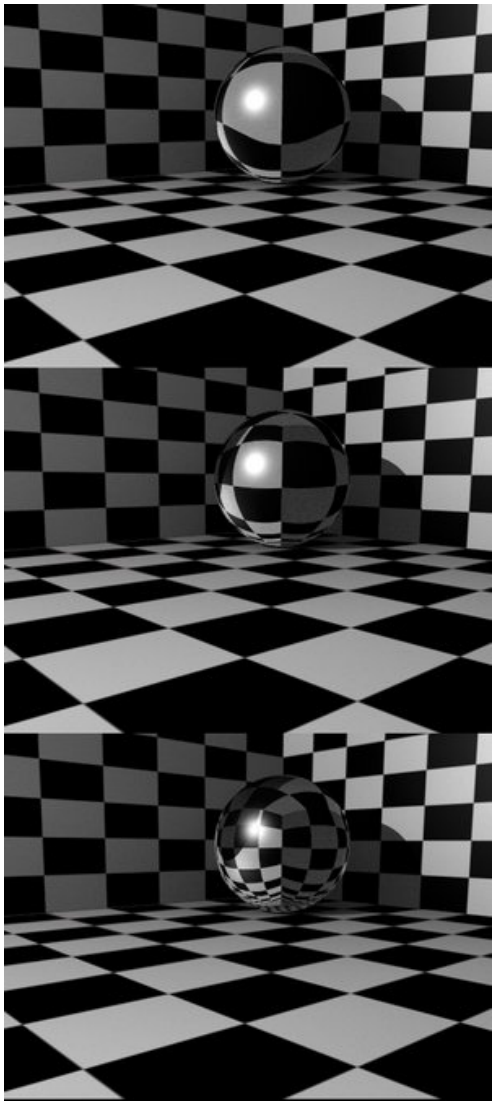
Fresnel
: Sets the power of the Fresnel effect. The Fresnel effect controls how transparent the Material is, depending on the angle between the surface normal and the viewing direction. Typically, the larger the angle, the more opaque a Material becomes (this generally occurs on the outline of the object).

Fac
: A controlling 'factor' to adjust how the blending (between transparent and non-transparent areas) happens.

SpecTra
: This slider controls the Alpha/falloff for Specular colors

Examples

Index of Refraction



Influence of the IOR of an Object on the distortion of the background: spheres of Water, Glass and Diamond (top to bottom).
(*Influence of the IOR of an Object on the distortion of the background: spheres of Water, Glass and Diamond (top to bottom).*). There are different values for typical materials: Air is **1.000** (no refraction), Alcohol is **1.329**, Glass is **1.517**, Plastic is **1.460**, Water is **1.333** and Diamond is **2.417**.

Glass

Realistic Glass is not perfectly transparent, and so light can reflect off of it, it can bend light, and you can even sometimes see shadows cast on the glass, since in reality the non-shadowed surface of the glass is reflecting some light back at you.



Realistic Glass Settings

In the Render panel, enable Raytracing and Shadows. In the example render shown, we have a lamp just off camera to the right, a red-colored ball, an angled piece of glass covering the bottom of the picture, a teal-colored NURBS torus (donut), and a back wall, in Z order. You can see:

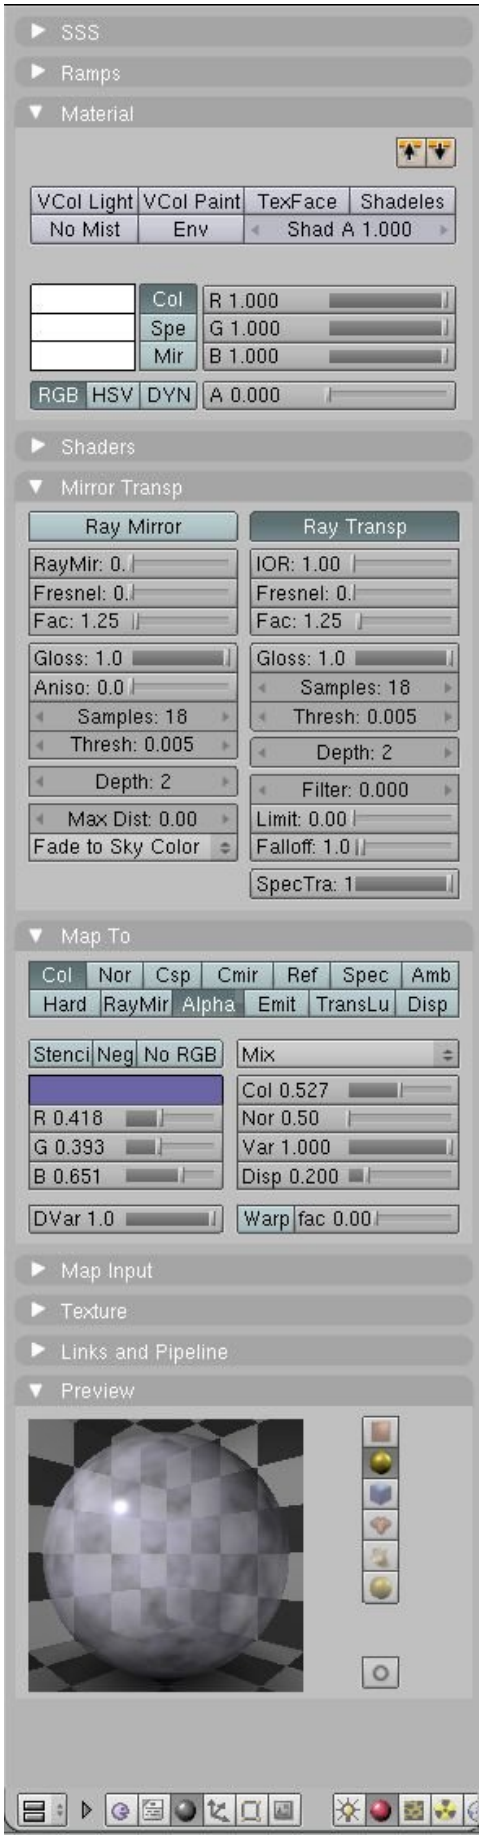
- Refracted wall, the donut, the environment behind the glass
- Shadows cast onto the glass by the red ball in front of the glass
- Shadows cast through the glass onto the donut by the red ball
- Shadow cast by the ball behind the glass into the box
- Refraction of the shadow of ball and torus coming back through the glass
- Specular glare on the glass from the lamp

In the Material settings, shown here for slightly blue green glass, enable Traceable and Shadbuff so that Raytraced as well as Shadow Buffer lamps will work with the glass.

In the Material panel, notice that the Alpha is not pure 0.0 In the Shaders panel, Shadow and TraShadow are enabled. The glass is colored by ambient light and emits its color, to simulate light bouncing around inside the glass itself and emerging

In the Mirror/Transp panel, default settings are used except for IOR.

Atmosphere/Cloud Cover



Cloud Cover Settings

In this example, we need a cloud cover for a planet. This is used to texture a UV Sphere that is a larger diameter than the planet inside of it.

The Alpha of the base material is zero, as is any Ambient or Emit values.

Of course, we have to use the Cloud texture, and put it to two uses: to slightly affect the color of the mesh, and to affect the Alpha transparency. While the base material is white, the texture gives it a blue tinge. It receives shadows, so the shadow of the planet will actually darken the dark side of the atmosphere.

Where the texture is black, think 0 color. If you map that 0 to alpha, you get transparent, or see-through. Where the texture is white, think full, 1.0 color. If you map white to alpha, you get full opaqueness. If that 1.0 also maps to some color, like blue, you will get a fully opaque blue color. In the example to the right, that full blue is then mixed down by the Col slider in the texture Map To panel so that it mixes with the base white material and only shades the white a slight hue of blue.

To make the clouds more pronounced, you can map the input to a higher value, and/or multiply (not mix) two texture channels layers together.

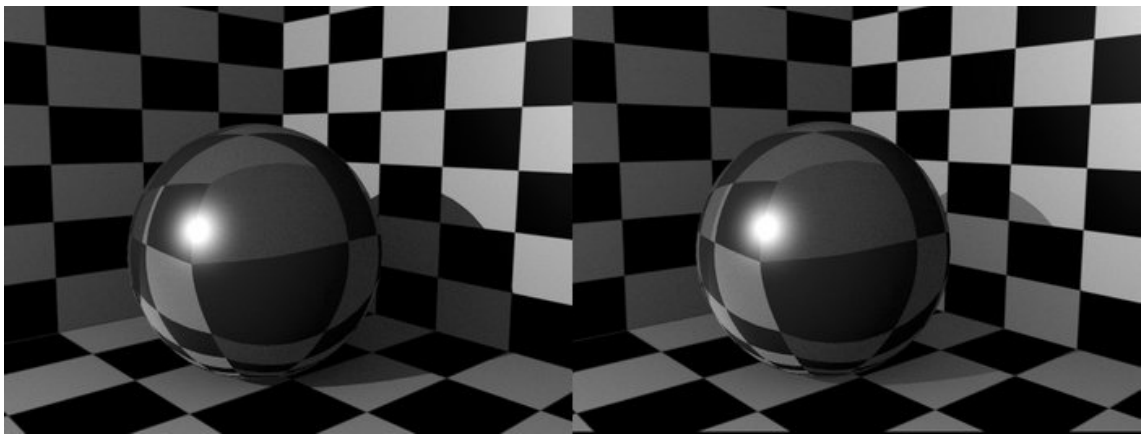
Hints

In order to get a physically accurate Fresnel effect with the current algorithm, you have to set Fresnel to 5.0 and Fac to 1.25. Nevertheless, you can play with these values for the sake of artistic freedom, if you feel the need to.

Casting Transparent Shadows

By default, the shadows of transparent objects are rendered solid black, as if the object was not transparent at all. But in reality, the more transparent an object is, the lighter its shadow will be. This can be taken into account, not in the Mirror Transp panel transparent object settings.

Transparent shadows are set on the materials that **receive** the shadows from the transparent object. This is enabled and disabled with the TraShadow button, in the Shading/Material context → Shaders Panel. The shadow's brightness is dependent on the Alpha value of the shadow casting Material.



Casting *transparent* shadows: TraShado 'off' on the left, TraShado 'on' to the right.

Ray versus Z Transparency



ZTransp, located in the Material Links panel, causes Blender to use the Z value, or distance from camera, to bleed things through. When using transparent planes with images, ray tracing will let the colors come through the transparent areas, but the Z values will not carry through. Therefore, later on down the rendering pipeline some issues may arise. For example, when OverSampling/Aliasing (OSA) is applied, artifacts may ensue as shown. In this example, raytracing is not providing the correct samples for the portion of the ground plane behind the transparent portions of the alpha-mapped tree image plane. To solve this, use ZTransp for the material, not raytracing.

As another alternative, change the texture interpolation filter size to 0.100 or some number smaller than 1.0. This slider is found in the Texture subcontext (F6 Map Image panel.

Raytracing affects the **Color**, but **does not affect the Alpha** channel value that is saved or passed on. Even if the Material is set to Alpha 0, raytracing will treat the material as transparent, but will return an alpha 1 for that channel regardless of the alpha setting inthe material.

ZTransp does affect Alpha channel values in the output. If you are going to be compositing the image later using Alpha over, you probably want to set ZTransp for your material. The disadvantage with ZTransp is that the light can not be bent or refracted.

To set the Alpha value of an object AND use raytracing, set the Pass Index the object (Object and Links panel). Render the image using Ray Transp and the Index pass (Render Layer panel). Then use the Index OB node to mask it out and the Set Alpha node to set the alpha for that portion of the image.

IOR values for Common Materials

The following list provides some Index Of Refraction values to use when RayTraced Transparency is used for various liquids, solids (gems), and gases:

A		E		J		S	
Acetone	1.36	Ebonite	1.66	Jade, Jadeite	1.64 - 1.667	Sanidine	1.522
Actinolite	1.618	Ekanite	1.600	Jade, Nephrite	1.600 - 1.641	Sapphire	1.757 - 1.779
Agalmatoite	1.550	Elaeolite	1.532	Jadeite	1.665	Sapphire, Star	1.760 - 1.773
Agate	1.544	Emerald	1.560 - 1.605	Jasper	1.540	Scapolite	1.540
Agate	1.540	Emerald Catseye	1.560 - 1.605	Jet	1.660	Scapolite, Yellow	1.555
Air	1.000		1.561	**K**			
Alcohol	1.329						
	1.36						

Alcohol, Ethyl (grain)	
Alexandrite	1.745
Alexandrite	1.750
Almandine	1.83
Aluminum	1.44
Amber	1.545
Amblygonite	1.611
Amethyst	1.540
Ammolite	1.600
Anatase	2.490
Andalusite	1.640
Anhydrite	1.571
Apatite	1.632
Apophyllite	1.536
Aquamarine	1.575
Aragonite	1.530
Argon	1.000281
Asphalt	1.635
Axenite	1.674 - 1.704
Axinite	1.675
Azurite	1.730
B	
Barite	1.636
Barytocalcite	1.684
Beer	1.345
Benitoite	1.757
Benzene	1.501
Beryl	1.57 - 1.60
Beryl, Red	1.570 - 1.598
Beryllonite	1.553
Brazilianite	1.603
Bromine (liq)	1.661
Bronze	1.18
Brownite	1.567
C	
Calcite	1.486
Calspar	1.486
Cancrinite	1.491
Carbon Dioxide (gas)	1.000449
Carbon Disulfide	1.628
	1.460
Emerald, Synth flux	
Emerald, Synth hydro	1.568
Enstatite	1.663
Epidote	1.733
Ethanol	1.36
Ethyl Alcohol	1.36
Euclase	1.652
F	
Fabulite	2.409
Feldspar, Adventurine	1.532
Feldspar, Albite	1.525
Feldspar, Amazonite	1.525
Feldspar, Labradorite	1.565
Feldspar, Microcline	1.525
Feldspar, Oligoclase	1.539
Flourite	1.434
Formica	1.47
G	
Garnet, Andradite	1.88 - 1.94
Garnet, Demantiod	1.880 - 1.9
Garnet, Demantoid	1.880
Garnet, Grossular	1.738
Garnet, Hessonite	1.745
Garnet, Mandarin	1.790 - 1.8
Garnet, Pyrope	1.73 - 1.76
Garnet, Rhodolite	1.740 - 1.770
Garnet, Rhodolite	1.760
Garnet, Spessartite	1.810
Garnet, Tsavorite	1.739 - 1.744
Kornerupine	1.665
Kunzite	1.660 - 1.676
Kyanite	1.715
L	
Labradorite	1.560 - 1.572
Lapis Gem	1.500
Lapis Lazuli	1.50 - 1.55
Lazulite	1.615
Lead	2.01
Leucite	1.509
M	
Magnesite	1.515
Malachite	1.655
Meerschaum	1.530
Mercury (liq)	1.62
Methanol	1.329
Milk	1.35
Moldavite	1.500
Moonstone	1.518 - 1.526
Moonstone, Adularia	1.525
Moonstone, Albite	1.535
Morganite	1.585 - 1.594
N	
Natrolite	1.480
Nephrite	1.600
Nitrogen (gas)	1.000297
Nitrogen (liq)	1.2053
Nylon	1.53
O	
Obsidian	1.489
Oil of Wintergreen	1.536
Oil, Clove	1.535
Oil, Lemon	1.481
Oil, Neroli	1.482
Oil, Orange	1.473
Oil, Safflower	1.466
Oil, vegetable (50Â° C)	1.47
Olivine	1.670
Scheelite	1.920
Selenium, Amorphous	2.92
Serpentine	1.560
Shampoo	1.362
Shell	1.530
Silicon	4.24
Sillimanite	1.658
Silver	0.18
Sinhalite	1.699
Smaragdite	1.608
Smithsonite	1.621
Sodalite	1.483
Sodium Chloride	1.544
Spessarite	1.79 - 1.81
Sphalerite	2.368
Sphene	1.885
Spinel	1.712 - 1.717
Spinel, Blue	1.712 - 1.747
Spinel, Red	1.708 - 1.735
Spodumene	1.650
Star Ruby	1.76 - 1.773
Staurolite	1.739
Steatite	1.539
Steel	2.50
Stichtite	1.520
Strontium Titanate	2.410
Styrofoam	1.595
Sugar Solution 30%	1.38
Sugar Solution 80%	1.49
Sulphur	1.960
Synthetic Spinel	1.730
T	
Taaffeite	1.720
Tantalite	2.240
Tanzanite	1.690-1.7
Teflon	1.35
Thomsonite	1.530

Carbon Tetrachloride	
Carbonated Beverages	1.34 - 1.356
Cassiterite	1.997
Celestite	1.622
Cerussite	1.804
Ceylanite	1.770
Chalcedony	1.544 - 1.553
Chalk	1.510
Chalybite	1.630
Chlorine (gas)	1.000768
Chlorine (liq)	1.385
Chrome Green	2.4
Chrome Red	2.42
Chrome Tourmaline,	1.61 - 1.64
Chrome Yellow	2.31
Chromium	2.97
Chrysoberyl	1.745
Chrysocolla	1.500
Chrysoprase	1.534
Citrine	1.532 - 1.554
Citrine	1.550
Clinohumite	1.625 - 1.675
Clinozoisite	1.724
Cobalt Blue	1.74
Cobalt Green	1.97
Cobalt Violet	1.71
Colemanite	1.586
Copper	1.10
Copper Oxide	2.705
Coral	1.486
Coral	1.486 - 1.658
Cordierite	1.540
Corundum	1.766
Cranberry Juice (25%)	1.351
Crocoite	2.310
Crysoberyl, Catseye	1.746 - 1.755
Crystal	2.000
Garnet, Uvarovite	1.74 - 1.87
Gaylussite	1.517
Glass	1.51714
Glass, Albite	1.4890
Glass, Crown	1.520
Glass, Crown, Zinc	1.517
Glass, Flint, Dense	1.66
Glass, Flint, Heaviest	1.89
Glass, Flint, Heavy	1.65548
Glass, Flint, Lanthanum	1.80
Glass, Flint, Light	1.58038
Glass, Flint, Medium	1.62725
Glycerine	1.473
Gold	0.47
H	
Hambergite	1.559
Hauyn	1.490 - 1.505
Hauynite	1.502
Helium	1.000036
Hematite	2.940
Hemimorphite	1.614
Hiddenite	1.655
Honey, 13% water content	1.504
Honey, 17% water content	1.494
Honey, 21% water content	1.484
Howlite	1.586
Hydrogen (gas)	1.000140
Hydrogen (liq)	1.0974
Hypersthene	1.670
I	
Ice	1.309
Idocrase	1.713
Iodine Crystal	3.34
Iolite	1.522 - 1.578
Iron	1.51
Onyx	1.486
Opal, Black	1.440 - 1.460
Opal, Fire	1.430 - 1.460
Opal, White	1.440 - 1.460
Oregon Sunstone	1.560 - 1.572
Oxygen (gas)	1.000276
Oxygen (liq)	1.221
P	
Padparadja	1.760 - 1.773
Painite	1.787
Pearl	1.530
Periclase	1.740
Peridot	1.635 - 1.690
Peristerite	1.525
Petalite	1.502
Phenakite	1.650
Phosgenite	2.117
Plastic	1.460
Plexiglas	1.50
Polystyrene	1.55
Prase	1.540
Prasiolite	1.540
Prehnite	1.610
Proustite	2.790
Purpurite	1.840
Pyrite	1.810
Pyrope	1.740
Q	
Quartz	1.544 - 1.553
Quartz, Fused	1.45843
R	
Rhodizite	1.690
Rhodochrisite	1.600
Rhodonite	1.735
Rock Salt	1.544
Rubber, Natural	1.5191
Ruby	1.757 - 1.779
Rum, White	1.361
Tiger eye	1.544
Topaz	1.607 - 1.627
Topaz, Blue	1.610
Topaz, Imperial	1.605 - 1.640
Topaz, Pink	1.620
Topaz, White	1.630
Topaz, Yellow	1.620
Tourmaline	1.603 - 1.655
Tourmaline	1.624
Tourmaline, Blue	1.61 - 1.64
Tourmaline, Catseye	1.61 - 1.64
Tourmaline, Green	1.61 - 1.64
Tourmaline, Paraiba	1.61 - 1.65
Tourmaline, Red	1.61 - 1.64
Tremolite	1.600
Tugtupite	1.496
Turpentine	1.472
Turquoise	1.610
U	
Ulexite	1.490
Uvarovite	1.870
V-W	
Wardite	1.590
Variscite	1.550
Water (0Â° C)	1.33346
Water (100Â° C)	1.31766
Water (20Â° C)	1.33283
Water (gas)	1.000261
Water 35'C (Room temp)	1.33157
Whisky	1.356
Willemite	1.690
Witherite	1.532
Vivianite	1.580
Vodka	1.363
Wulfenite	2.300
Z	

Cuprite	2.850	Ivory	1.540	Rutile	2.62	Zincite	2.010
D						Zircon	1.777 - 1.987
Danburite	1.627 - 1.641					Zircon, High	1.960
Danburite	1.633					Zircon, Low	1.800
Diamond	2.417					Zirconia, Cubic	2.173 - 2.21
Diopside	1.680						
Dolomite	1.503						
Dumortierite	1.686						



Image 1a: Marble Dog with SSS. Watch especially the ears and the paws.



Image 1b: And the same without SSS.

Many organic and some inorganic skins are not totally opaque right at the surface, so light does not just bounce off the top surface. Instead, some light also penetrates the skin surface, and scatters around inside, taking on the color of the insides and emerging back out to blend with the surface reflection. Human/animal skin, the skin of grapes, tomatoes, fruits, wax, gels (like honey, or Jello) and so on all have subsurface scattering (SSS), and photo-realism really cannot be achieved without it.

SSS can be found in the Material buttons (F5), and is limited to diffuse shading only, it does not affect specular shading.

How it works



Image 2: First pass of SSS.
Actually calculating the light path beneath the surface of an object would be practically impossible. But it has been shown that it is not necessary to do this, and that one can use a different approach.

Blender calculates SSS in two steps:

- At first the brightness of the surface is calculated, from the frontside of the object as well as from it's backside. This is pretty much the same as in a normal render. Ambient-Occlusion, Radiosity, the type of diffuse Shader, the light color etc. is taken into account (*Image 2*).
- In the second step the image is rendered finally, but now the SSS shader replaces the diffuse shader. Instead of the lamps the calculated lightmap is used. The brightness of a surface point is the calculated "Average" of the brightness of it's surrounding points. Depending on your settings the whole surface may be taken into account, and it's a bit more complicated than simply calculating the average, but don't bother too much with the math behind it.

Instead let's see what SSS does to a distinct light point.

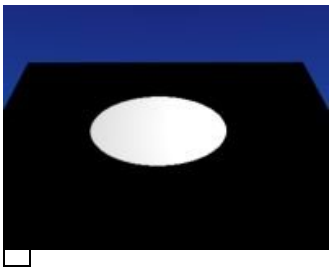


Image 3a: No SSS.

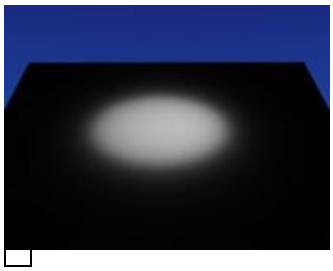


Image 3b: Small SSS radius.

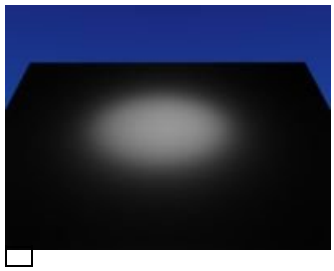


Image 3c: SSS radius enlarged.

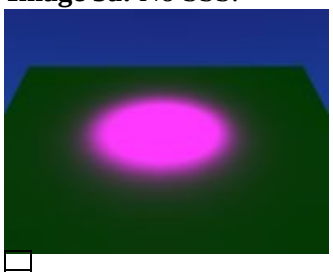


Image 3d: SSS with very large green radius value.
If you turn on SSS the light is distributed over a larger Area. The size of this area depends on the radius values. Instead of distributing all colors with the same amount, you may choose different radius values for each of the RGB-Colors.

If you use a very large radius value for a color, its light is evenly distributed over the whole object.

Enabling SubSurface Scattering

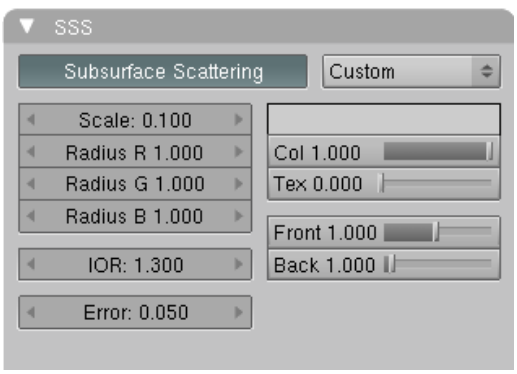


Image 4: The SSS Panel. SSS is already enabled.

- Enable SSS by clicking on the Subsurface Scattering button. You will see your preview panel render change somewhat, as additional processing kicks in.
- Various pre-sets are defined for you, selected by clicking the Selector to the right of Custom. If you don't like any of them, you can define a Custom set. When you select a pre-set, the *Radius* values, the color and the *IOR* are set for you. The remaining options are not set (because they are mostly dependent on the size of your object).

SubSurface Scattering doesn't need raytracing. But since it is dependent on the incident light and shadows, you need proper shadow calculation (which may need raytracing).

Options

The numeric sliders control how the light is scattered:

Scale
: The scale of your object, in Blender units, across which you want the scattering effect to take place. For the presets scale **1.0** means **1** Blender unit equals **1** millimeter, scale **0.001** means **1** Blender unit equals **1** meter. If you want to work out what scale value to use in your scene, just use the formula: (size in blender units)/(real world size in millimetres)= scale

Radius R, Radius G and Radius B
: The light blurring radius. As the light travels through the object and back up to emerge from the surface at some other point, it creates a path length. These sliders allow you to adjust the average length of that path. The longer the path length is, the more evenly this color is distributed.

IOR
: The IOR value determines the falloff of incident light. Higher values means that light falls off faster. The effect is quite subtle and changes the distribution function only a little bit. By the examination of many different materials a value of **1.3** to **1.5** have been found fitting. If you know the exact material you are trying to simulate, see our IOR table.

Error
: This parameter controls how precisely the algorithm samples the surrounding points. Leaving it at **0.05** should give images without artifacts. It can be set higher to speed up rendering, potentially with errors. Setting it at **1.0** is a good way to quickly get a preview of the look, with errors.

The color swatch and blend control the color of the SSS shader.

<swatch>
: This has two effects:
 1. If you think of the SSS as a strange sort of lamp, this would be the lights color.
 2. It also affects the scattering - the darker the color the more the light is scattered.

 So if you set it to green, the lit areas of the object will appear in green, and green is scattered only a little. Therefore the darker areas will appear in red and blue. You can compensate the different scattering by setting a larger radius for the color.

Col
: This controls how much the R, G, B option modulates the diffuse color and textures. Note that even with this option set to **0.0**, the R, G, B option still influences the scattering behavior.

Tex
: How much the surface texture is blurred along with the shading.

Front
: Factor to increase or decrease the frontscattering. When light enters through the front of the object, how much is absorbed or added? (Normally **1.0** or **100%**).

Back
: Factor to increase or decrease the backscattering. Light hitting an object from behind can go all the way through the object and come out on the front of the object. This happens mostly on thin objects, like hands and ears.

Developing your own SSS material

Follow these simple steps to make your own SSS material:

- Set the SSS color on a value of your choice, normally the predominant color of the object. If you want to use different radiuses for the colors, don't make it too dark.
- Set the scale factor. If you want to see much translucency you need small objects or large scale values.
- Set the radius values.
- Adjust the brightness with the Front and Back values.

Example: Grapes

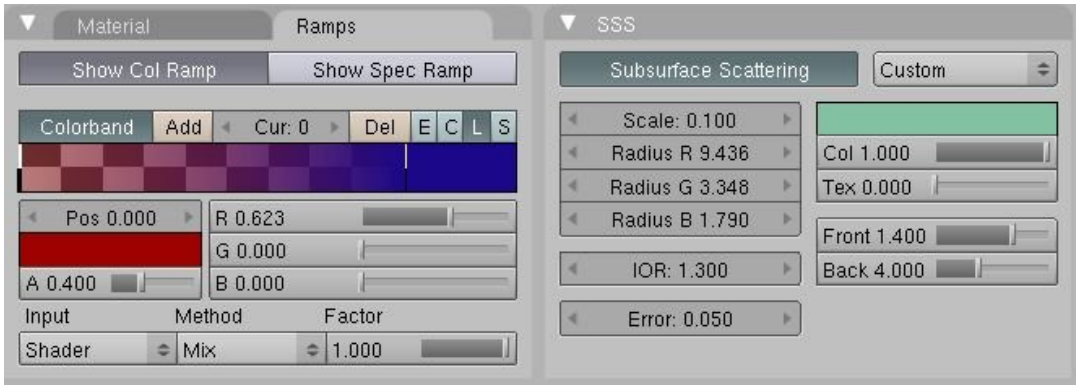
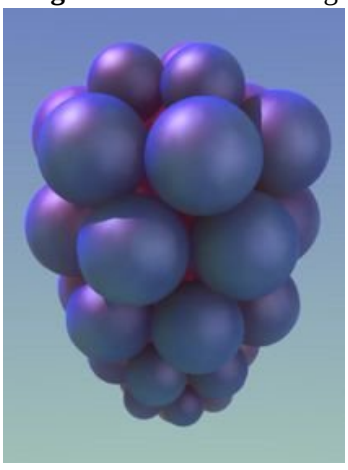


Image 5: Subsurface Settings for the Grapes in Image 6.





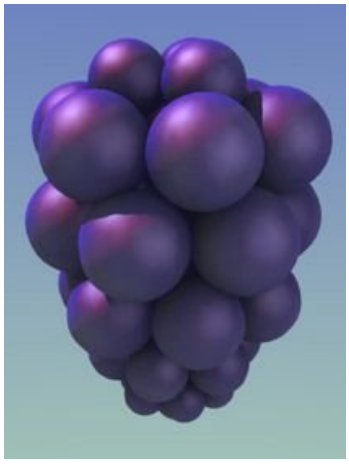




Image 6a: With SSS.

Image 6b: Difference between 6a and 6c, with strongly enhanced brightness and contrast.



Image 6c: Without SSS.

The skin of the grape is a purple colorramp, and we observe that grapes have a fairly red specular glow. The scene is lit with a bright sun from above and behind, and a wide soft area light as a key light. A cloud texture is used to introduce surface variations.

In the example in (*Image 6a*), we have SSS turned on to give a green color based on the inside of a grape. The red *Radius*-Value is quite large, the green *Radius*-Value is larger than the blue one. We can observe the effects of these settings in (*Image 6b*). Though the SSS color is green, the green values are only increased at the very bright spots on the grapes. Green and blue are nearly equally scattered (the larger radius for green compensates the green SSS-color). Since red is scattered very much, red is missing on the parts that are lit from front. The red light is scattered all over each grape, so the same amount of light is emitted from a larger area, partially from the backside of the grapes.

Where we see the backsides of the grapes (pointing away from the light) they appear in red. This has two reasons:

1. Red is scattered stronger than green and blue, so more of the red light reaches the backside of the grapes.
2. The *Back*-Light setting is strongly increased. The *Front*-Light setting is slightly raised to compensate the loss in brightness from the scattering.

Get the .blend file

Example: Skin



Skin is the holy grail of materials, because it is so varied, so imperfect, and so complex. A good skin render is a combination of procedural, UV-mapped images for color, normal, specularity, ambient, and so on. This example uses SSS to get you started.

The model was a human **1.75 BU** high (each BU=1m in real world). We wanted a Caucasian human, so we started with a light tan base material, very little hardness and specularity. For SSS, we started with the "Skin 1" preset. The head was **0.25 BU** in diameter, hence the SSS Scale of **0.150**, because we do not want light from one side to light up the other; there is supposed to be a skull in there!

Lighting plays an important part in getting the basic skin to look right. For this example I used a 3-point studio rig:

- Key: Spot placed **5 BU** from subject. Energy **2.0**, Falloff **5.0**, color (**0.98**, **0.98**, **1.0**).
- Fill lights: Hemi placed **5 BU** out to the side, **1 BU** in front of subject. Energy **0.5**, Falloff **10**, color white.

Strands

Strands are the rendered pathvisualisations of a particle animation. There are two different strand methods available:

- **Polygonstrands:** This is the default (old) method. The strands are rendered as flat polygons. The number of polygons depend on the Render steps settings in the Visualizations panel of the Object context, Particle sub-context.
- **Keypointstrands:** You activate Keypointstrands with the button Strand render in the Visualization panel of the particle system. The hair curves are not stored as polygons, but only the key points, which are then converted to polygons on the fly. A second difference is the way transparency works. Rather than rendering them using the existing system, all strand segments in a part are sorted front to back and rendered in that order.

Keypointstrands

◊ Are more memory efficient and faster, to make rendering of large amounts of fur and grass possible. For good performance, the render steps button should be lowered (e.g. 2 should be good enough fur), since the result will be a smoothed curve anyway. You need 1 to 2 render steps less than steps in the 3D window. Also, using more render parts helps to reduce memory usage.
◊ Have a distance of vision reduction (Simplification panel) for children from faces.
◊ May be faded out towards the tip without an additional texture.
◊ Are not raytraced. So they are not visible through raytransparent materials or in a raymirror (you can use *Environment Mapping* for that).
◊ Have shape problems if they are rendered with a greater width.
◊ Can not carry an UV-Texture along the strand.

Polygonstrands

◊ Work well with greater width, so you can use them as an alternative to billboards because the strands may have an animated shape.
◊ Can be textured with an UV-Texture along the strands.
◊ Are seen by raytracing.

Strands Shading

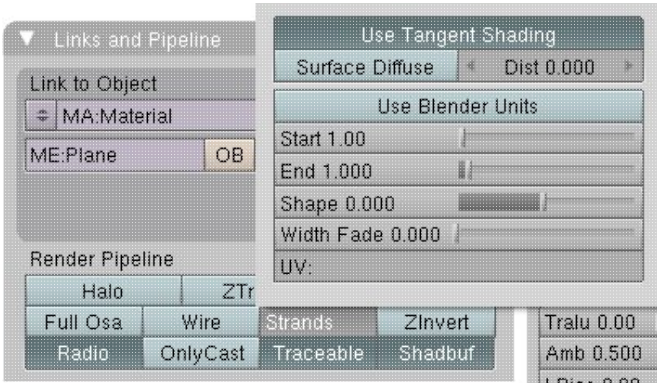


Image 1: Strands shader settings in the Links and Pipeline panel.
Strands are rendered with the material of the underlying face/vertex, including shading with an UV-Texture. Since you can assign more than one material to each face, each particle system may have itâ← s own material and the material of the underlying face can be different from the material of the strands.

Additionally strands can be shaded along the strand (from root to tip) with a mono-dimensional texture, only polygonstrands can carry a two-dimensional UV-Texture.

The options for the strand shading are hidden behind the Strands button in the Links and Pipeline panel of the Shading context, Material sub-context.

Use Tangent Shading
: Calculates the light as if the strands were very thin and round. This makes the hair appear brighter and shinier. Disabling the â← Tangent Shadingâ← option will still render nicely, but resembles more solid strands, as though made of metal or wood.

Surface Diffuse
: Computes the strand normal taking the normal at the surface into account. This eases the coloring and lighting of hair a lot, especially for keypointstrands. Essentially hair reacts similar like ordinary surfaces and donâ← t show exaggerated strong and large specular highlights.

Dist
: The distance in Blender units over which to blend in the normal at the surface (if you want to use Surface Diffuse only for Grass/Fur in greater distance).

Use Blender Units
: Normally strands are quite thin, the thickness is given in screenpixels. If you use Blender units (BU) you may set the start value up to 2 BU, the end value up to 1 BU. You have to consider the overall object size, because the smallest possible size is 0.001 BU. So if you use 1 BU for 1 meter the smallest possible size would be 1 mm (too thick for thin hair).

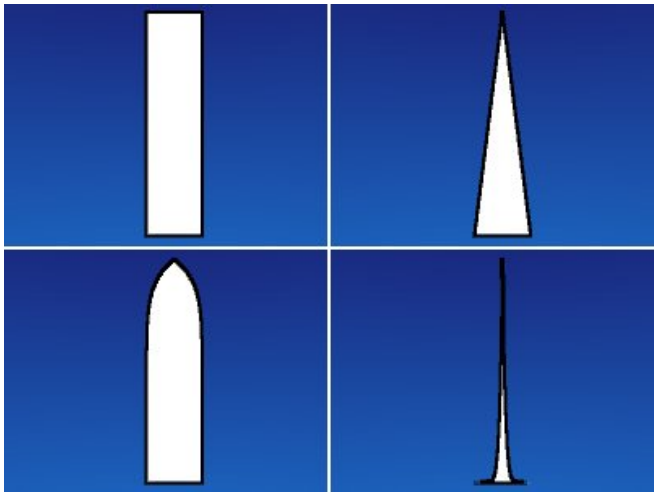


Image 2: a) Sta=End, b) End=0.0, Shape=0.0, c) Shape=0.9, d) Shape=-0.9.

Start
: Width of the hair at the root.

End
: Width of the hair at the tip.

Shape
: This slider allows you to control the interpolation. Default (0.0) is a linear interpolation between Start and End. A negative value will make the strand narrower (spiky), a positive value will make it fatter.

Width Fade

To fade out along the width of the strand. This works only for keypointstrands. 0.0 is no fading at all, 1.0 linear fading out.

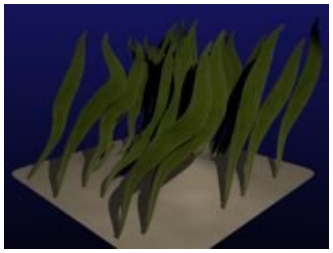


Image 3: Sea weed something.

UV

You can texture the polygonstrands with an UV-Texture. Fill in the name of the UV-Set (not the texture) here. You have to load the texture also in the Shading context, Texture and Material sub-contexts (Map Input: UV; you may use every Map To setting you like - especially the alpha value, see *Image 3*).

Texturing along the Strand



Image 4: Fading a strand to alphaâ←¦



Image 5: â←¦And the render result.

Strands can be textured along the strand, i.e. from root to tip. To do that you have to activate the Strand button in the Map Input panel of the Material sub-context.

Pretty much the most important setting is shown in (*Image 4*), how to fade the tip of a strand to alpha to make nice, fuzzy looking hair. Normally you would use a linear blendtexture for this.

You may of course set any attribute you like, especially color. Be careful with specularity, hairs tend to get to shiny.

Strand render Simplification

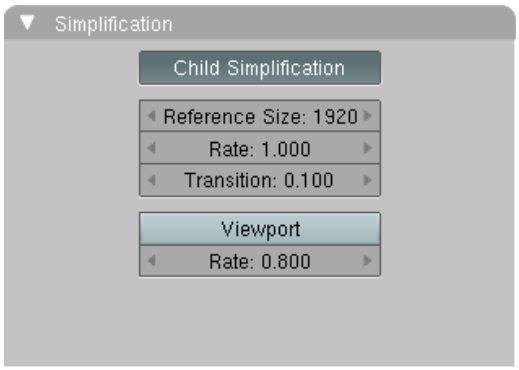


Image 5: Strand render child simplification.
If you use keypointstrands (Strand render button) and have activated Children from: Faces, the panel Simplification appears. The strand render has options to remove child strands as the objectâ← s faces become smaller.

Reference Size
: This is the approximate size of the object on screen (in pixels), after which simplification starts.

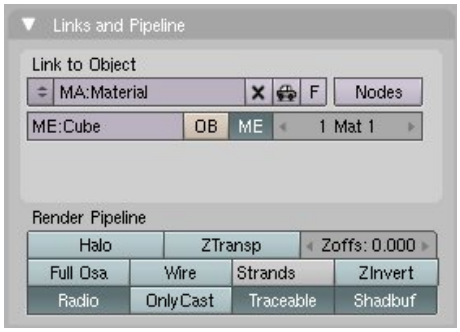
Rate
: How fast strands are removed.

Transition
: The transition period for fading out strands as they are removed.

Viewport
: This removes strands on faces that are outside of the viewport.
 Rate
 : Controls how fast these are removed.

Introduction

In addition to creating materials as just described using all the settings on all the materials panels, Blender allows you to create a material by routing basic materials through a set of nodes. Each node performs some operation on the material, changing how it will appear when applied to the mesh, and passes it on to the next node. In this way, very complex material appearances can be achieved.



Links and Pipeline tab

You should already be familiar with general material concepts and how to create materials/textures using the material panel. You should also have a general understanding of the texture coordinate systems available in Blender (e.g. Orco, UV, etc.). Also, when reading this I intend to purposely skip aspects of a node because in later sections you will see the function expanded upon. Each section builds off the previous.

I will begin by saying that the node system does not make the material pane obsolete. Many features and material settings are still only accessible through the material panel (e.g. Ray Mirror). However with the advent of nodes, more complex and fantastic materials can be created since we now have greater control.

Just in case you're not (yet) familiar with the concepts: when you create a system of nodes (otherwise known as a "noodle"), you're describing a data-processing pipeline of sorts, where data "flows from" nodes which describe various *sources*, "flows through" nodes which represent various processing and filtering stages, and finally "flows into" nodes which represent outputs or destinations. You can connect the nodes to one another in many different ways, and you can adjust "knobs," or parameters, that control the behavior of each node. This gives you a tremendous amount of creative control. And, it will very quickly become intuitive.

Having said all that, letâ← s begin with a normal material (*Links and Pipeline tab*).

Here we have the standard material we have added to a cube mesh. I could, as I have in the past, add color and other settings to this material and it would certainly look nice. But letâ← s say I am just not getting what I am looking for? What if I what to control the creation more tightly or add more complexity? Here is where nodes come in.

Making this node map is accomplished by working in a Node Editor window. This section covers:

- Node editor window and basic controls
- How to work with a node (general)
- The specific types of nodes available for materials

Accessing The Node Editor

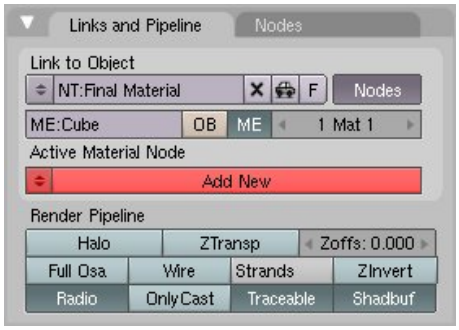


Select the Node Editor window
First lets enter the node editor (*Select the Node Editor window*) and make sure that the node editor has the material node button (the sphere icon) pressed (*Node Editor Toolbar*), not the composite node button.

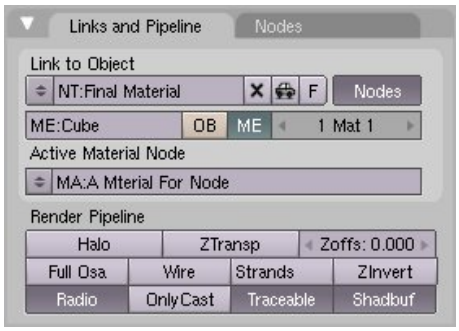


Node Editor Toolbar for Materials

Enabling Node Materials in the Material Buttons



Links and Pipeline Node tab
Letâ← s take the base material (*Links and Pipeline tab*) and hit the Nodes button next to the material name in the material panel or the node editor. You will see a change in the material panel (*Links and Pipeline Node tab*).



New Node Tree Material
What you have just done is told Blender to make the material you were on (in this case "Final Material") to become the node tree (hence the new name - NT: versus MA:). Under the node tree you can see that it is asking you to add a new material (*Links and Pipeline Node tab*). Once you do (*New Node Tree Material*) you will create a material (MA:) that is under node tree. After adding a material to the node tree, two nodes will appear in the node editor - a material node, and an output node.

It is important to note that you can add a new material (which you can edit and change like any other material in the material panel), add an already created material or append a material from another blender file, and also use the material that you used to create the node tree.

External Links

- Material Nodes Overview - Explains the Material node editor, material node types and some examples of node setups and node-groups.
- Blender Material Nodes - Changelog for the Blender version that introduced material nodes.

This section explains the window in general, and its header menu options. It also tells you how to enable nodes for use within Blender.

More information on **where to use nodes** please refer to the following pages:

- Nodes for Materials (types of material nodes).
- Nodes for Composition (types of composition nodes).

To learn more about **how to handle nodes themselves** in general please refer to the following page:

- Doc:Manual/Materials/Nodes/Usage.

Also there is a reference page available explaining the nodes windows and connected functions:

- Doc:Reference/Windows/Nodes.

Accessing The Node Editor



Select the Node Editor window.
First let's enter the node editor by changing our window type to Node Editor. As shown in *Select the Node Editor window*, click on the window type icon and select Node Editor from the popup list. Node maps can get quite large, so use or create a big window. The window has a graph-paper style background and a header.

Each scene within your blend file can have multiple Material Node map and ONE Compositing Node map. The Node Editor window shows either map, depending on the selector position.

Hint
You might want to add a new window layout called 6-Nodes (the list is shown on the User Preferences header at the top of your screen) comprised mostly of one big Node Editor window. My layout has the buttons window at the bottom and a text editor window on the side for me to keep notes. If you have a widescreen display (or even a regular one), you might also want to add a 3D view or UV/Image Editor window to the left side of the Node window layout, so you can work with images or your model while you're manipulating nodes. Having the 3D Preview Render panel open on top of an object is quite useful if you're tweaking material nodes.

By default, the header, when first displayed, is uninitialized as shown:



Default Node Editor header.

Activating Nodes

- What nodes to use?
 - ♦ If you want to work with a material node map, click the ball in the Material/Compositing node set selector. (see *Node Editor Header with Material Nodes enabled.*)
 - ♦ If you want to work with a compositing node map, click the face on the Material/Compositing node set selector. (see *Node Editor Header with Compositing Nodes enabled.*)
- To actually activate nodes, click the Use Nodes button.
- The first time that you select either a Material or a Compsiting node map, the Node Editor window will be instantly filled with starter input and output compositing nodes already connected together.

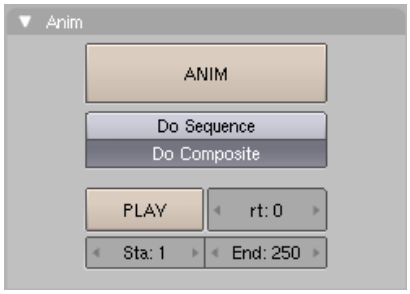


Node Editor Header with Material Nodes enabled.



Node Editor Header with Compositing Nodes enabled.

Required settings for Composition



Use Composition Nodes.
If you are compositing, you must now tell Blender to use the Node map that has been created, and to composite the image using the Node Map. To do so, click on the Do Composite button located below the Animation button. This tells Blender to composite the final image by running it through the composition node map.

From here, you add and connect nodes in a sort of map layout to your heart's content (or physical memory constraints, whichever comes first). But first, let's lay the groundwork by going over the window in general and the header menu options and buttons.

Node Editor Window Actions

When the cursor is in the window, several standard Blender hotkeys and mouse actions are available, including:

Popup menu
: Space - Brings up a main popup menu, allowing you to add, view, select, etc.

Delete
: X or Del - Deletes the selected node(s).

Box select
: B - Starts the bounding box selection process. Position your cursor and LMB click & drag to select a set of nodes.

Cut connections (box)
: LMB click & drag - Starts a box selection, BUT when you let up the mouse button, all threads (connections) within the box are broken.

Undo
: CtrlZ Very helpful if you forgot to press B before box-selecting, eh?

Redo
: CtrlY or Ctrlâ § ShiftZ - You can use this if you used "undo" a bit to often :)

Select multiple
: â § ShiftLMB or â § ShiftRMB - Multiple node select.

Grab/Move
: G - Moves your current selection around.

Execute
: E - pumps inputs through the noodle, refreshing everything.

Standard Window Control
Node maps can get pretty hairy (large and complicated, that is). The contents of the window, (the node map) can be panned just like any other Blender window by clicking MMB and dragging about. Wheeling Wheel up/down or using the keypad
+ NumPad/- NumPad will zoom in/out. The window can be resized and combined using the standard window techniques (see *Navigating in 3d Space*).

Node Editor Header



Node Editor Header with Material Nodes enabled.



Node Editor Header with Compositing Nodes enabled.

On the window header, you will see header options:

- View - to see things more clearly;
- Select - to do things more clearly;
- Add - to walk with...err..to add Nodes, organized by type;
- Node - to do things with selected nodes, akin to vertices;
- a Material or Compositing node set selector;
- a Use Nodes button;
- a Free Unused button.

View, Select, and Add Header Menus

These popup menus provide the basic functions:

View
: This menu changes your view of the window, standing in for the standard keyboard shortcuts + NumPad (zoom in), - NumPad (zoom out), â Home (zoom all) or equivalent mouse actions.

Select
: This menu allows you to select a node or groups of nodes, and does the same as typing the hotkey to select all A or start the border select B process.

Add
: This menu allows you to add nodes. Please see the next section for a discussion on the types of nodes that you can add, and what they do. Clicking this menu item is the same as pressing Space when the cursor is in the window

Node Header Menu

Show Cyclic Dependencies
: C - Ok, so you've been adding and connecting nodes to your heart's content, and you haven't run out of memory yet. Selecting Show Cyclic Dependencies will show you where you have connected your threads in a circle. For example, you can easily connect a mix output as input to another node, and then connect that node's output back to the mix node input, resulting in a little circle where the image just runs round and round. Left alone, it will eventually get tired and dizzy and crash your computer.

Hide
: H - Hides your selected nodes. Just like vertices in a mesh.

Grouping
: Most importantly, this menu option allows you to create a user-defined group of nodes. This group can then be edited and added to the map. To create a group, select the nodes you want, and then Node â Make Group, or just use the keyboard shortcut CtrlG. Edit the name using the little input box in the group. Groups are easily identified by their green header and cool names you have picked for them.

Delete
: X - Deletes selected nodes.

Duplicate

â § ShiftD - Makes an Unlinked copy, with the same settings as the original.

Grab
G - Moves the little nodes around according to your mouse, just like with meshes.

Duplicate - Faked you out
The new copy is placed **exactly over the old one**. But it isn't the connected one, so playing with the controls will do nothing to your images, even though it **looks** like it's connected with the little threads coming out of the node that is **underneath**. You have to move the duplicated node to reveal the connected node beneath it.

Grab - Reminder Only
Just like my mother-in-law, the menu item does not actually do anything; it's just there to remind you that you can press the G key when your cursor is in the window and actually accomplish something with your life (like rearranging nodes in the window).

Material/Composting Selector Button

Nodes are grouped into two categories, based on what they operate on. Material Nodes operate on a material in use within the blend file. To work with Material Nodes, click on the ball. When you want to work with Compositing nodes, click on the face to show the Compositing Node map.

Use Nodes Header Button

This button tells the render engine to use the node map in computing the material color or rendering the final image, or not. If not, the map is ignored and the basic render of the material tabs or scene is accomplished.

Free Unused Header Button

This button frees up memory space when you have a very complex node map. Recommended.

Please only proceed after reading Node Editor. You need to be in the Node editor window in order to follow the text below.

Node Basics

What are nodes

"Nodes" (some people call connected nodes "noodles") are individual blocks that perform a certain operation. Conceptually, there are three basic types of nodes:

- Nodes that *produce* information, but do not have any inputs of their own. The **Render Layers**,**Value** and **RGB** nodes are examples of this.
- Nodes that *filter* or *transform* their inputs, to produce one or more outputs. The **RGB Curves, Defocus,** and **Vector Blur** nodes are an example.
- Nodes that *consume* their inputs to produce some kind of meaningful result, such as **Composite** (which determines the final output used by Blender), **Viewer** (which displays things right in the

noodle display, and **File Output**. A single noodle can produce many different outputs.

The essential idea of nodes is that you can create an arbitrarily-complex *network* of nodes, by connecting the *outputs* of one or more nodes to the *inputs* of one or more other nodes. Then, you can set appropriate parameters (as you see fit) for each node. This network ("noodle") describes how image and texture information literally *flows through* the "noodle" to produce whatever result you want.

As you'll see in just a moment, you can also define *node groups*, then treat those groups as though they were a single node. (Rather like a "macro" or a "user-defined function" in a programming language...) You can link and append these node groups from other files.

Adding and Arranging Nodes

Nodes are added in two ways to the node editor window:

- By clicking the Add menu in the node editor toolbar and picking the type of node you want, or
- By having your cursor in the node editor window and pressing Space and picking a node from the popup Add menu.

In general, try to arrange your nodes within the window such that the image flows from left to right, top to bottom. Move a node by clicking on a benign area and dragging it around. The node can be clicked almost anywhere and dragged about; connections will reshape as a bezier curve as best as possible.

Threads beneath Nodes
Threads (the curves that connect sockets) may reposition behind a node; however they are just that and do not interact with that node in any way.

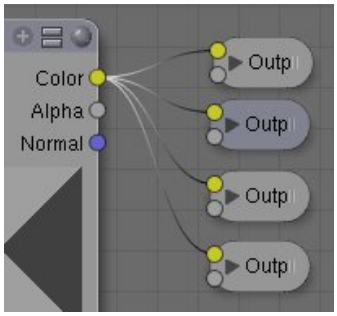
Nodes beneath Nodes
"Gotcha!" When you duplicate a node, the new node will be positioned *exactly* on top of the node that was duplicated. If you leave it there (and it's quite easy to do so), you can't easily tell that there are *two* nodes there! Sure, the drop-shadow may be slightly darker than before, but it's no longer visually obvious just what's connected to what. When in doubt, grab a node and move it slightly to see if something's lurking underneath.

Sockets



Node Sockets.
Each Node in your node window will have "sockets" (often also referred to as "connectors") which are small colored circles to which input data and output data will be linked (*Node Sockets*).

The sockets on the left side of a node describe *inputs*, while the sockets on the right side are *outputs*.



Node Linking.
For your convenience, nodes are *color-coded* according to the type of information they expect to send or receive. There are three colors:

Yellow sockets
: Indicates that **color** information needs to be input or will be output from the node.

Grey sockets
: Indicates values (**numeric**) information. It can either be a single numerical value or a so-called "value map." (You can think if a value map as a grayscale-map where the different amount of bright/dark reflects the value for each point.) If a single value is used as an input for a "value map" socket, all points of the map are set to this same value.
Common use: Alpha maps and value-options for a node.

Blue/Purple sockets
: Indicates **vector/coordinate/normal** information.

Between nodes, yellow must be linked to yellow, gray to gray, blue to blue, unless you use a *converter*, which we'll cover later on.

Next to the color in the node you will see the name of that socket. Though not always the case, you can see the name of the socket as what the information is *intended* to be. But this is not necessarily what it *has* to be. For example, I can add a link from an gray socket titled Alpha to the material node's gray Reflection socket and still get a result, they key thing being that it's a "gray to gray" connection.

There are exceptions where you can mix yellow (*e.g.* a color-image) and gray (*e.g.* grayscale) without converters. Blender normally places a converter if needed, so feel free to experiment with them. You can use the "Viewer" output nodes, as explained in the later sections, to see if/how it works.

Connecting and Disconnecting Sockets

You link between sockets by clicking the socket with the LMB and holding to drag the thread to another socket, you then let go once you reach the corresponding socket. To break a link between sockets click the LMB and hold and drag a box around any part (it can be really small) to break the link. From output sockets, multiple threads can be extracted and attached to many nodes (*Node Linking*).

In this case, a copy of each output is routed along a thread. However, only a single thread can be linked to an input socket.

Node Groups

Both material and composite nodes can be grouped. Grouping nodes can simplify the node network layout in

the node editor, making your material or composite 'noodle' (node network) easier to work with. Grouping nodes also creates what are called NodeGroups (inside a .blend file) or NodeTrees (when appending).

Conceptually, "grouping" allows you to specify a *set* of nodes that you can treat as though it were "just one node." You can then re-use it one or more times in this or some different .blend file(s).

As an example: If you have created a material using nodes that you would like to use in another .blend file, you *could* simply append the material from one .blend file to another. However, what if you would like to create a new material, and use a branch from an existing material node network? You could re-create the branch. Or you could append the material to the new .blend file, then cut and paste the branch that you want into the new material. Both of these options work, but are not very efficient when working across different .blend files. What if you have created a â← Depth of Fieldâ← composite node network and would like to use it in another .blend file? What if you wanted to apply exactly the same series of operations, dozens of times? Here again, you *could* re-create the network, but this is not very efficient. A better method of re-use, for either material node branches or composite node networks, would be to create groups of nodes.

Once a group has been defined, it becomes an opaque object; a reusable software component. You can (if you choose) ignore exactly how it is *defined*, and simply use it (as many times as you like) for whatever it *does.* Groups can be made available through the Blender's library and standard appending method.

Grouping Nodes

To create a node group, in the node editor, select the nodes you want to include, then press CtrlG or Space â node â make group. A node group will have a green title bar. All of the selected nodes will now be minimized and contained within the group node. Default naming for the node groups is *NodeGroup, NodeGroup.001* etc. There is a name field in the node group you can click into to change the name of the group. Change the name of the node group to something meaningful. When appending node groups from one .blend file to another, Blender does not make a distinction between material node groups or composite node groups, so I recommend some naming convention that will allow you to easily distinguish between the two types. For example, name your material node branches *Mat_XXX*, and your composite node networks *Cmp_XXX*.

What NOT to include in your groups.

Remember that the essential idea is that a group should be an easily-reusable, self-contained, software component.

Material node groups should not include material nodes or output nodes.

If you include a material node in your group, you'll end up having the material node appear *twice:* once inside the group, and once outside the group in the new material node-network. If you include an output node *in* the group, there won't be an output socket available *from* the group!

Composite node groups cannot include a *Render Layer* node (Blender won't let you), and should not contain a *Composite* output node. Here again, if they include any connected Output node (*Viewer, Split Viewer*, etc.) then the Group will not have an image output socket.

Editing Node Groups

With a group node selected, pressing Tab expands the node to a window frame, and the individual nodes within it are shown to you. You can move them around, play with their individual controls, re-thread them internally, etc. just like you can if they were a normal part of your editor window. You will not be able to thread them to an outside node directly from them; you have to use the external sockets on the side of the Group node. To add or remove nodes from the group, you you need to ungroup them.

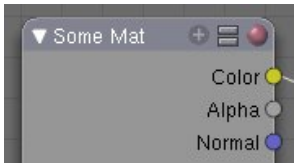
Ungrouping Nodes

The AltG command destroys the group and places the individual nodes into your editor workspace. No internal connections are lost, and now you can thread internal nodes to other nodes in your workspace.

Appending Node Groups

Once you have appended a NodeTree to your .blend file, you can make use of it in the node editor by pressing Space â Add â Groups, then select the appended group. The "control panel" of the Group is the individual controls for the grouped nodes. You can change them by working with the Group node like any other node.

Node Controls



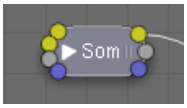
Top of a Node.
At the top of a node there are up to 4 visual controls for the node (*Top of a Node*). Clicking these controls influences how much information the node shows.

 **Arrow**
The arrow on the left collapses the node entirely (*Collapsing Arrow*).

 **Plus sign (+)**
The "plus" icon collapses all sockets that do not have a thread connected to it (*Plus Sign*).

 **Two squares (=)** or "Equal sign"
The icon with the two squares collapse all the items in a node that have boxes with information in them (*Menu Collapse*).

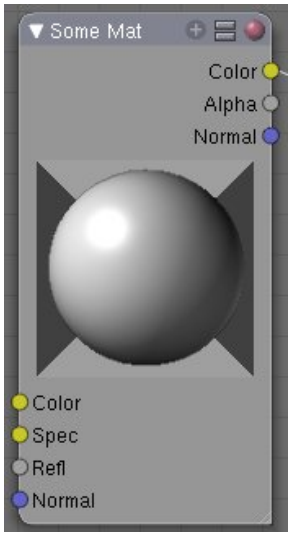
 **Sphere**
The sphere icon collapses the viewing window (if the node has one) (*Sphere*).
If the Sphere is **red** this can have 3 reasons:
- ◊ It's the only effective output Composite node in the compositor.
- ◊ It's the only effective material Output node (the first one that is added).
- ◊ If it's a Material input node that has a Material (MA:) assigned to it.



Collapsing Arrow.



Plus Sign.



Menu Collapse.

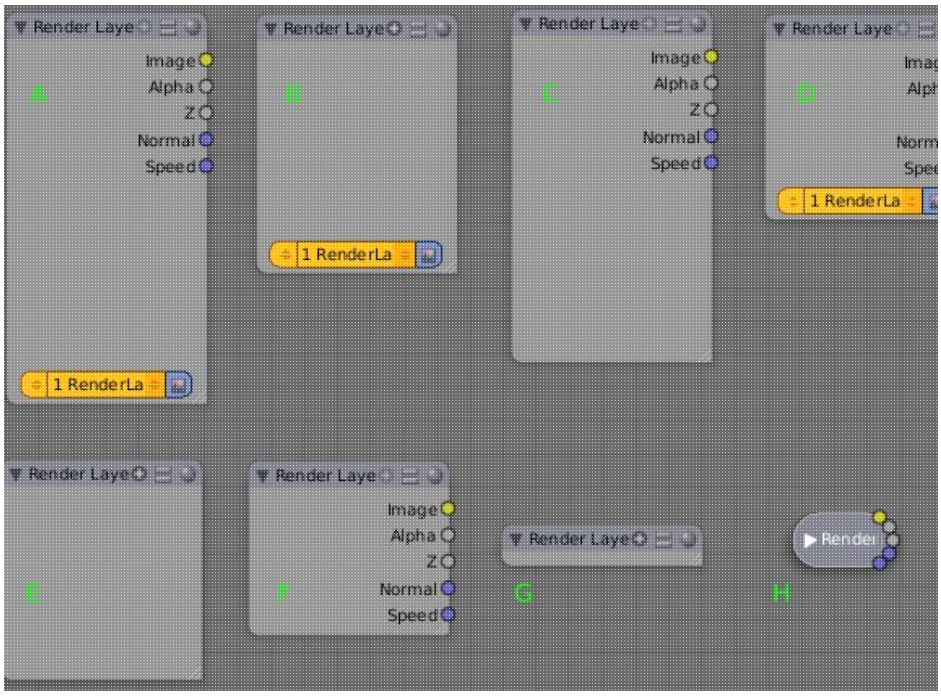


Sphere.



In Combination.

The later three can be used in varying combinations with each other. The arrow that collapses the entire node can only be used in combination with the plus sign (*In Combination*).



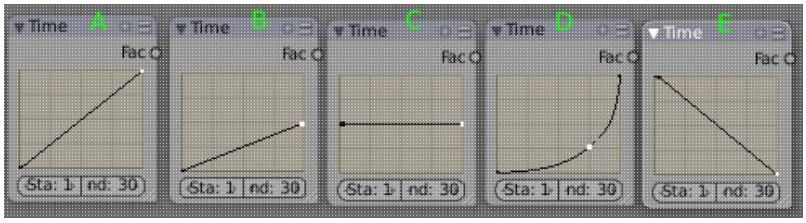
Top sizing controls of a Node
A) Normal, **B**) + Sign clicked, **C**) = Sign clicked, **D**) Sphere clicked, **E**) + and = clicked,
F) = and Sphere clicked, **G**) All three clicked **H**) Arrow clicked.

Node Sizing

Fine Sizing of an individual node can also be accomplished somewhat by clicking LMB and dragging in the lower right-hand corner (where the little slanted lines are).

Node Curves

Some nodes have a curve area that translates an input value to an output value. You can modify this curve shape by clicking on a control point and moving it, or adding a control point. Some examples are shown below:



Modifying a curve node.
Every curve starts out as a straight line with a slope of 1. (My daughter NEVER thought she would use her high school algebra. Ha!) The curve starts out with two tiny black control points at each end of the line. Clicking LMB on a control point selects it and it turns white.

Changing the curve affects how the output is generated. The input, X, usually proceeds linearly (at regular intervals) across the **bottom** axis. Go up until you hit the curve, and then over to the **right** to determine the Y output for that corresponding X. So, for the second example, as X goes from 0 to 1.0 across the bottom, Y varies from 0.0 to 0.5. In the third, as X goes from 0.0 to 1.0 across the bottom, Y stays constant at 0.5. So, in the picture above, these curves have the following affect on time: **A** don't affect, **B** slow down, **C** stop, **D** accelerate, and **E** reverse time.

The "Curves" widget is a built-in feature in Blender's UI, and can be used anywhere, provided the curve data itself is being delivered to this widget. Currently it is in use in the Node Editor and in the UV Window.

This widget will map an input value horizontally and return the new value as indicated by the height of the curve.

Note: The fact that one of the points on the curve is "white" in each of these screen-shots is *not* significant: it just means that it happened to be the point most-recently selected by your author when preparing this tutorial. What matters here is the shape of *the curve*, not the position (nor the color) of the control-points that were used to define it.

RGB Curves

Multiple curves can be edited in a single widget. The typical use, RGB curves, has "Combined" result or "Color" ("C") as the first curve, and provides curves for the individual R, G, and B components. All four curves are active together, the "C" curve gets evaluated first.

Selecting curve points

- LMB always selects 1 point and deselects the rest.
- Hold â § Shift while clicking to extend the selection or select fewer points.

Editing curves

- LMB click&drag on a point will move points.
- A LMB click on a curve will add a new point.
- Dragging a point exactly on top of another will merge them.
- Holding â § Shift while dragging snaps to grid units.
- CtrlLMB adds a point.
- Use the X icon to remove selected points.

Editing the view

The default view is locked to a 0.0-1.0 area. If clipping is set, which is default, you cannot zoom out or drag the view. Disable clipping with the icon resembling a #.

- LMB click&drag outside of curve moves the view
- Use the + and - icons to zoom in or out.

Special tools

The wrench icon gives a menu with choices to reset a view, to define interpolation of points, or to reset the curve.

Please be sure to read **Nodes** before reading this page.

The material nodes allow you to manipulate a material by routing it through a map of connected nodes. A starting material is routed through different nodes that do various things to the material, combined with other inputs or put back together, and finally output where it can be applied to a mesh, halo, particle, etc.

Materials can be split into their RGB components, combined (mixed) with other inputs, and layered on top of one another. In the real world, this process is accomplished by mixing specific paints, thinners, surface prep, and then painting using various techniques under environmental conditions.

This section is organized by type of node, which are grouped based on similar functions:

- **Input** - Introduces a material or component to the node map.
- **Output** - Displays the result in progress as a small image.
- **Color** - Manipulates the colors of the material.
- **Vector** - Change the way light is reflected off the material.
- **Convertors** - Convert colors to other material colors.
- **Groups** - User-defined groups of nodes.
- **Dynamic** - Custom nodes defined by Python. These are also known as PyNodes.

The simplest way to add a node is to put your cursor in a Node Editor window, and then press Space and click on Add. The popup menu will extend to show you these types of nodes. Click on a node type and the popup menu will extend again to show you specifc node types available.

Material Input Nodes

A starting material is created in the Materials Panel. The Nodes button is enabled to add that material to the list of noded materials shown in the Node Editor window header. Other inputs to the node map include:

- A value
- A color
- A texture
- Geometry

Material Node

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Input â Material



Material node
The Material node is used to add a material to the node program. Materials can be anything from pure shading to fully layered with textures. It inputs the main attributes of a material (color, alpha and normal vector) into the map.

Output

Materials can output color (which includes shading and any textures assigned to it), alpha, and the final normal calculated from any textures it has.

- Color
- Alpha
- Normal

Input

Materials can take inputs for colors, inputs for diffuse color and specularity color, a value for reflectivity, and a normal.

- Color - The base color of the paint. Can be set
 - ♦ manually by LMB clicking on the color swatch applet next to the socket, choosing a color using the control panel that pops up, and pressing â µ Enter
 - ♦ based on an Active Material which is specified using the material panels, or
 - ♦ plugged in from an RGB color generator.
- Spec - The color that is reflected as you get perpendicular to the light source reflecting off the surface. The color can be
 - ♦ plugged in from another node or
 - ♦ set manually by LMB clicking on and using the color swatch applet.
- Refl: - The degree to which the material reflects light and gives off its color. The value can be provided by another node or set manually.
- Normal - The lighting condition.

Controls

MA:Material field
: You can browse and select materials here.

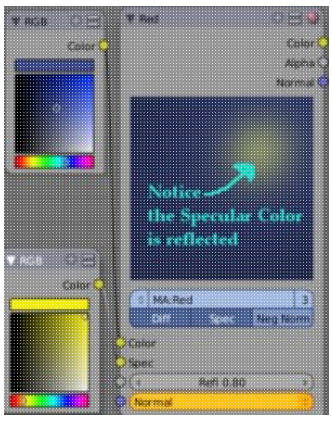
Diff toggle
: Turn on/off Diffuse Color.

Spec toggle
: Turns on/off Specularity calculation.

Neg toggle
: Inverts the material input normal when activated (which, of course, is a combination of the 3D normal given to it by the 3D object plus the normal input point).

Normal Override
The normal input socket does not in any way blend the source normal with the underlying geometry. Any plugged in Geometry here overrides the Normal lighting conditions.

Using the Material Node with Specularity



Material Node using Specularity

To make a material node actually generate a color, you have to specify at least a basic input color, and optionally a specularity color. The specularity color is the color that shines under intense light.

For example, consider the mini-map to the right. The base color, a dark blue, is connected from an RGB color generator node to the Color input socket. The specular color, yellow, is connected to the Spec input. Under Normal lighting conditions on a flat surface, this material will produce a deep blue color and, as you approach a spot perpendicular to the light, you will see the yellow specular color mix in.

Enable Spec
To see specularity, you have to enable it by clicking the blue Spec button located just below the material color swatch in the node.

Value Node

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Input â Value



Value Node
The Value node has no inputs; it just outputs a numerical value (floating point spanning 0.00 to 1.00) currently entered in the NumButton displayed in its controls selection.

Use this node to supply a constant, fixed value to other nodes' value or factor input sockets.

RGB Node

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Input â RGB



RGB Node
The RGB node has no inputs. It just outputs the Color currently selected in its controls section; a sample of it is shown in the top box. In the example to the right, a gray color with a tinge of red is slected.

To change the brightness and saturation of the color, LMB click anywhere within the square gradient. The current saturation is shown as a little circle within the gradient. To change the color itself, click anwhere along

the rainbow Color Ramp.

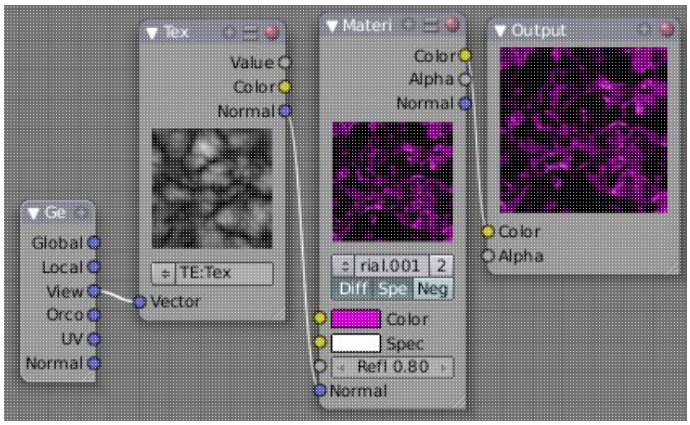
Texture Node

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Input â Texture

A texture, from the list of textures available in the current blend file, is selected and introduced through the value and/or color socket.

Library
Please read up on the Blender Library system for help on importing and linking to textures in other blender files



In the example to the right, a cloud texture, as it would appear to a viewer, is added to a base purple material, giving a velvet effect.
Note that you can have multiple texture input nodes. In the (old) panel way, multiple textures were assigned to channels and each channel was checked on or off to be applied to the material. With nodes, you simply add the textures to the map and plug them into the map.

Geometry Node

Panel: Node Editor → Material Nodes

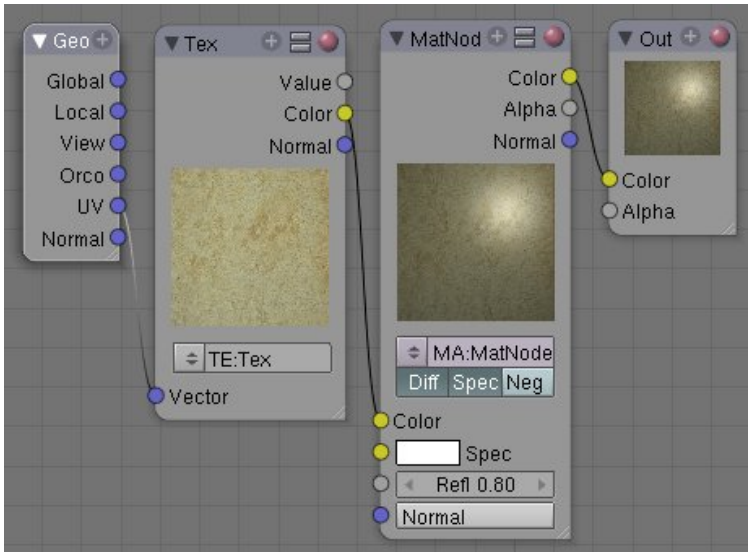
Menu: â § ShiftA â Input â Geometry



Geometry node
The geometry node is used to specify how light reflects off the surface. This node is used to change a material's Normal response to lighting conditions.

Note
These are exactly the same settings as in the Map Input panel for Textures, though a few settings - like Stress or Tangent - are missing here. Normally you would use this node as input for a Texture Node.

Geometry Node Example using a UV image

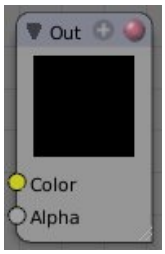


Setup to render an UV-Mapped Image Texture.
E.g.: To render an UV-mapped image, you would use the UV output and plug it into the Vector Input of a texture node. Then you plug the color output of the texture node into the color input of the material node - which corresponds to the setting on the Map To panel.

Material Output Nodes

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Output â Output



Output node
At any point, you may want to see the work in progress, especially right after some operation by a node. Simply create another thread from the output socket of the node to the picture input socket of an Output node to see a mini-picture.

Connect the alpha channel to set/see transparency.

Effective Output Node
The only Output node that is used for the Material in the end (i.e the only non-Preview) has a little **red sphere** on the upper right.

Material Color Nodes

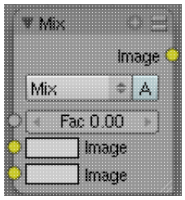
These nodes play with the colors in the material. The choices are:

- Mix
- RGB Curves

Mix Node

Panel: Node Editor â Material Nodes

Menu: â § ShiftA â Color â Mix



This node mixes a base color or image (threaded to the top socket) together with a second color or image (bottom socket) by working on the individual and corresponding pixels in the two images or surfaces. The way the output image is produced is selected in the drop-down menu. The size (output resolution) of the image produced by the mix node is the size of the base image. The alpha and Z channels (for compositing nodes) are mixed as well.

Not one, not two, but count 'em, *sixteen* mixing choices include:

Mix
: The background pixel is covered by the foreground using alpha values.

Add
: The pixels are added together. *Fac* controls how much of the second socket to add in. Gives a bright result.
The "opposite" to Subtract mode.

Subtract
: The foreground pixel (bottom socket) is subtracted from the background one. Gives a dark result.
The "opposite" to Add mode.

Multiply
: Returns a darker result than either pixel in most cases (except one of them equals white=**1.0**). Completely white layers do not change the background at all. Completely black layers give a black result.
The "opposite" to Screen mode.

Screen
: Both pixel values are inverted, multiplied by each other, the result is inverted again. This returns a brighter result than both input pixels in most cases (except one of them equals **0.0**). Completely black layers do not change the background at all (and vice versa) - completely white layers give a white result.
The "opposite" of Multiply mode.

Overlay

A combination of Screen and Multiply mode, depending on the base color.

Divide
: The background pixel (top socket) is divided by the second one: if this one is white (= **1.0**), the first one isn't changed; the darker the second one, the brighter is the result (division by **0.5** - median gray - is same as multiplication by **2.0**); if the second is black (= **0.0**, zero-division is impossible!), Blender doesn't modify the background pixel.

Difference
: Both pixels are subtracted from one another, the absolute value is taken. So the result shows the distance between both parameters, black stands for equal colors, white for opposite colors (one is black, the other white). The result looks a bit strange in many cases. This mode can be used to invert parts of the base image, and to compare two images (results in black if they are equal).

Darken
: Both pixels are compared to each other, the smaller one is taken. Completely white layers do not change the background at all, and completely black layers give a black result.

Lighten
: Both parameters are compared to each other, the larger one is taken. Completely black layers do not change the image at all and white layers give a white result.

Dodge
: Some kind of inverted Multiply mode (the multiplication is replaced by a division of the "inverse"). Results in lighter areas of the image.

Burn
: Some kind of inverted Screen mode (the multiplication is replaced by a division of the "inverse"). Results in darker images, since the image is *burned* onto the paper, er..image (showing my age).

Color
: Adds a color to a pixel, tinting the overall whole with the color. Use this to increase the tint of an image.

Value
: The RGB values of both pixels are converted to HSV values. The values of both pixels are blended, and the hue and saturation of the base image is combined with the blended value and converted back to RGB.

Saturation
: The RGB values of both pixels are converted to HSV values. The saturation of both pixels are blended, and the hue and value of the base image is combined with the blended saturation and converted back to RGB.

Hue
: The RGB values of both pixels are converted to HSV values. The hue of both pixels are blended, and the value and saturation of the base image is combined with the blended hue and converted back to RGB.

Color Channels

There are (at least!) two ways to express the channels that are combined to result in a color: RGB or HSV. RGB stands for the Red,Green,Blue pixel format, and HSV stands for Hue,Saturation,Value pixel format.

A

Click the green Alpha button to make the mix node use the Alpha (transparency) values of the second (bottom) node. If enabled, the resulting image will have an Alpha channel that reflects both images' channels. Otherwise, (when not enabled, light green) the output image will have the Alpha channel of the base (top input socket) image.

Fac

The amount of mixing of the bottom socket is selected by the Factor input field (Fac:). A factor of zero does not use the bottom socket, whereas a value of **1.0** makes full use. In Mix mode, **50:50** (**0.50**) is an even mix between the two, but in Add mode, **0.50** means that only half of the second socket's influence will be applied.

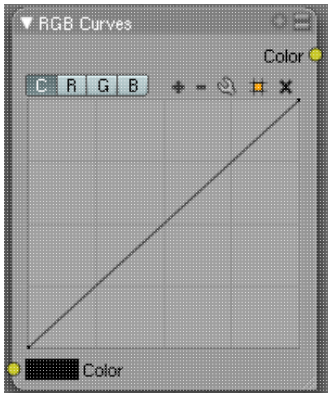
Using Dodge and Burn (History Lesson)

Use the dodge and burn mix methods in combination with a mask to affect only certain areas of the image. In the old darkroom days, when yes, I actually spent hours in a small stinky room bathed in soft red light, I used a circle cut out taped to a straw to dodge areas of the photo as the exposure was made, casting a shadow on the plate and thus limiting the light to a certain area. To do the opposite, I would burn in an image by holding a mask over the image. The mask had a hole in it, letting light through and thus 'burning' in the image onto the paper. The same equivalent can be used here by mixing an alpha mask image with your image using a dodge mixer to lighten an area of your photo. Remember that black is zero (no) effect, and white is one (full) effect. And by the way, ya grew to like the smell of the fixer, and with a little soft music in the background and the sound of the running water, it was very relaxing. I kinda miss those dayz.

RBG Curves Node

Panel: Node Editor → Material Nodes

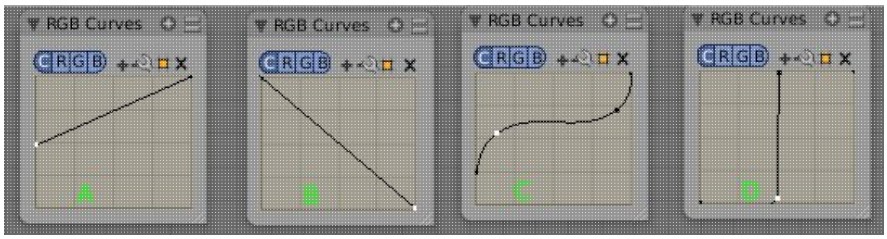
Menu: â § ShiftA â Color â RGB Curves



RGB Curves node
For each color component channel (RGB) or the composite (C), this node allows you to define a bezier curve that varies the input (across the bottom, or x-axis) to produce an output value (the y-axis). By default, it is a straight line with a constant slope, so that .5 along the x-axis results in a .5 y-axis output. Click and drag along the curve to create a control point and to change the curve's shape. Use the X to delete the selected (white) point.

Clicking on each C R G B component displays the curve for that channel. For example, making the composite curve flatter (by clicking and dragging the left-hand point of the curve up) means that a little amount of color will result in a lot more color (a higher Y value). Effectively, this bolsters the faint details while reducing overall contrast. You can also set a curve just for the red, and for example, set the curve so that a little red does not show at all, but a lot of red does.

Here are some common curves you can use to achieve desired effects:



A) Lighten **B**) Negative **C**) Decrease Contrast **D**) Posterize

Material Vector Nodes

Mapping Node

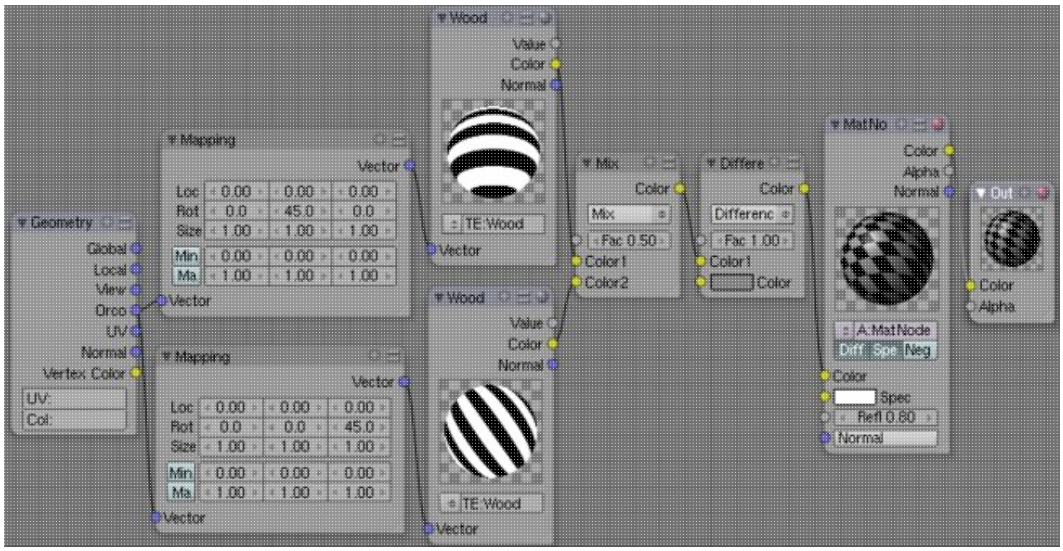
Panel: Node Editor â Material Nodes

Menu: â § ShiftA â Vector â Mapping

Essentially mapping node allows the user to modify a mapping. It is possible to use it do same operations as the Map Input panel found in Material buttons allows. It also makes it possible to do several things that are not possible in Map Input. Mapping can be rotated and clamped if desired. Currently mapping node supports only flat mapping type though.

The controls of the node have been ordered in X, Y, Z order. If you want to use the clamping options, try enabling Min and Max.

Mapping Node Example

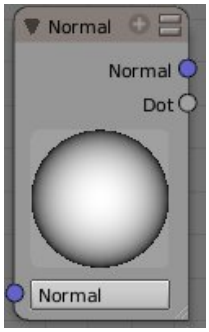


Using mapping nodes to produce amazing chess checkers texture.
This simple example shows one possible way to use mapping nodes to produce interesting textures. As you can see simply by mapping same texture a bit differently can make it useful. Using the same technique it would be easy to mimic different kind of stripy fabrics.

Normal Node

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Vector â Normal



Normal node
The Normal node generates a normal vector and a dot product. Click and Drag on the sphere to set the direction of the normal.

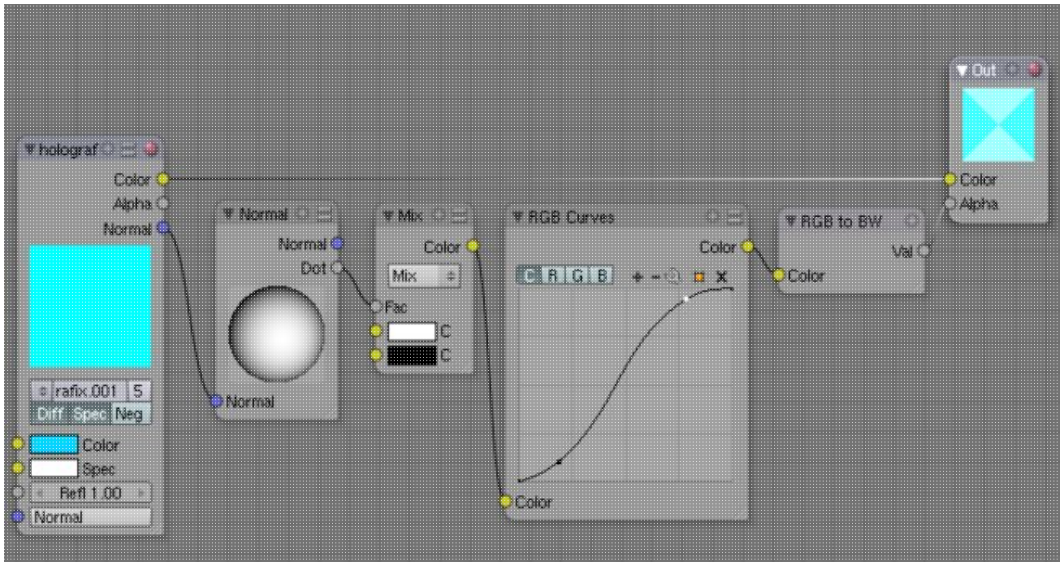
This node can be used to input a new normal vector into the mix. For example, use this node as an input to a Color Mix node. Use an Image input as the other input to the Mixer. The resulting colorized output can be easily varied by moving the light source (click and dragging the sphere).
The (face) normal is the direction of the face in relation to the camera. You can use it to do the following:

- Use this node to create a fixed direction -> output Normal.
- Calcuate the Dot-Product with the Normal-Input. The Dot-Product is a scalar value (a number).
 - ♦ If two normals are pointing in the same direction the Dot-Product is 1.
 - ♦ If they are perpendicular the Dot-Product is zero (0).
 - ♦ If they are antiparallel (facing directly away from each other) the Dot-Product is -1. *And you never thought you would use the Vector Calculus class you took in college - shame on you!*

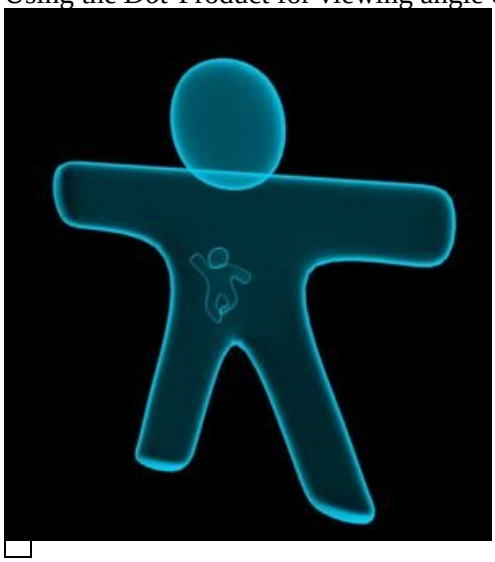
So now we can do all sorts of things that depends on the viewing angle (like electron scanning microscope effect or some of techniques described in the section Doc:Tutorials/Textures/Map_Input_Techniques). And the best thing about it is that you can manipulate the direction interactively.

One caveat
The normal is evaluated per face, not per pixel. So you need enough faces, or else you don't get a smooth result

Normal Node Example



Using the *Dot*-Product for viewing angle dependent material, in this case the *Alpha*-Value.

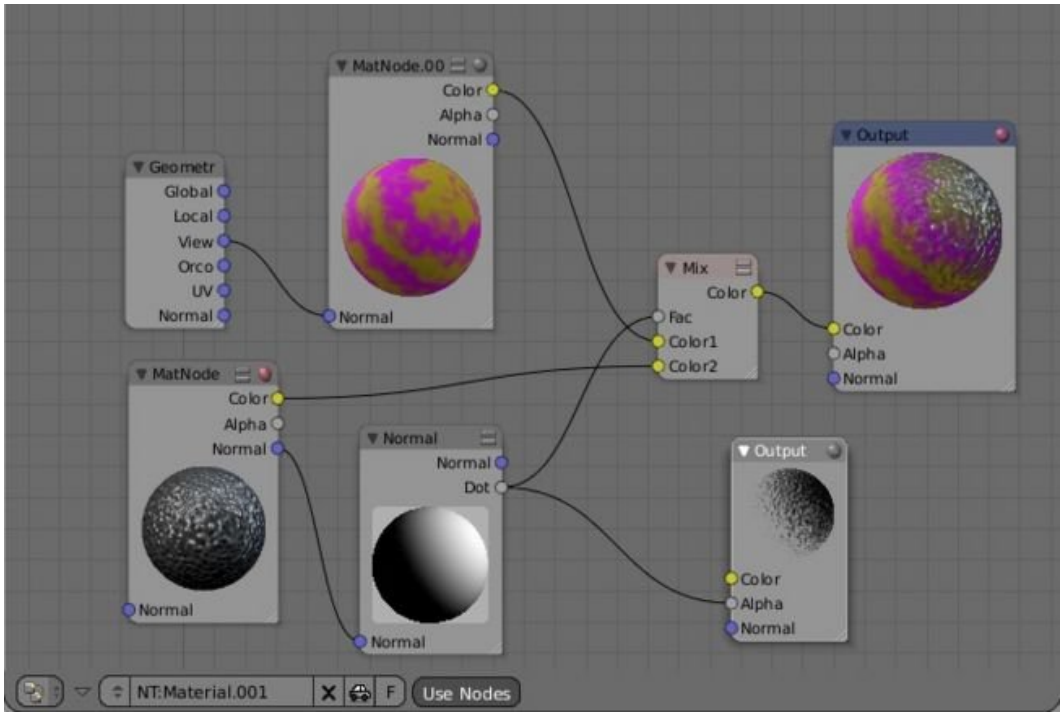


What a surprise - Gus is pregnant!
In the shown example the Dot-Product is used to govern the Alpha-Value of the material. The RGB Curves node is used to sharpen the blend between black and white (which results in a range from fully transparent to fully opaque).

The material is an ordinary blue/cyan material with a high Emit value and Z-Transp activated; the background is black. So, as the face angle gets closer to pointing directly at the camera, the material gets more transparent. As the face gets closer to pointing at a right angle to the camera (in this case facing up or down), it gets more opaque.

Using Normal to reflect bumpy-looking textures

We can use the Normal node to shift the reflection, and thus the shinyness, of a material as shown below. This effect can also be done without nodes, using the materials and texture panels. However, using Nodes allows you to graphically see what is going on to create the final material. The result is that, even though the surface mesh in your model is physically smooth, the reflection makes it look like it has a very fine bumpyness or uneven coating to it, like a really bad paint job or surface prep job was done. This is key to making realistic-looking surfaces.

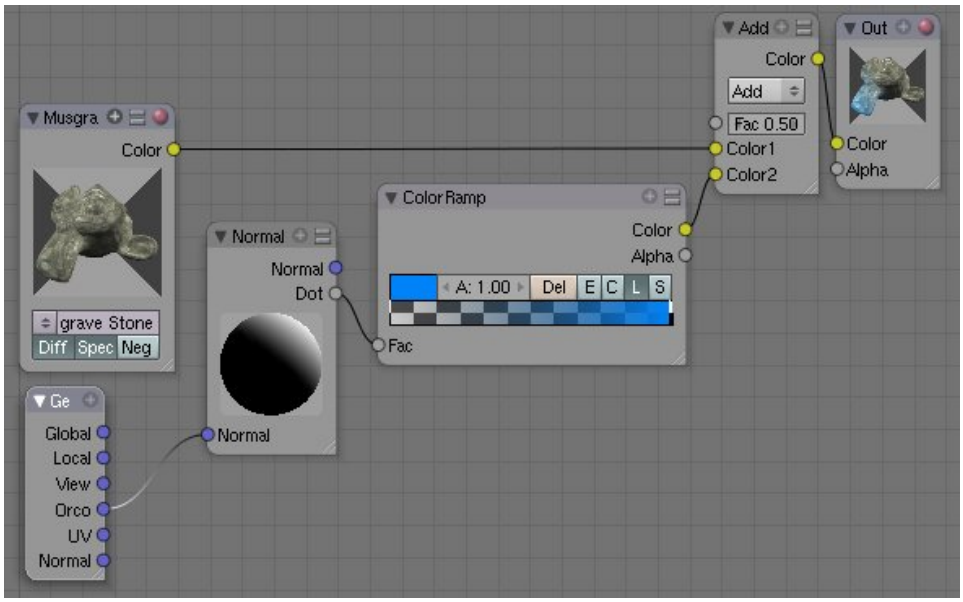


Using Normal for bump effect

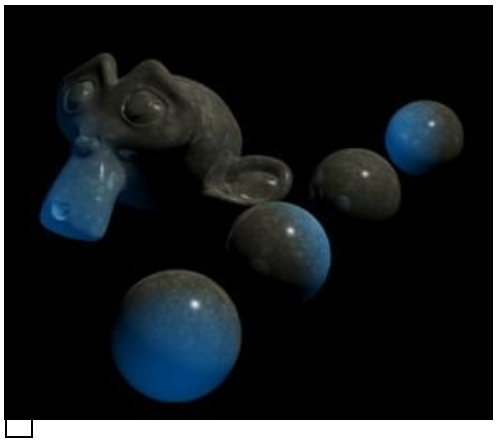
This map starts out using the geometry node to modify a base material called MatNode.001, shown in the header as MatNode.00 (the "1" was cut off to make the Node skinny, which may happen to you). The Geometry node puts out the View vector set, which flattens the color (removes any shine). MatNode.001 is itself a base tan color with a marble texture that affects and mixes in a purple where the texture is black.

A second material, MatNode, which is the default gray color with a noise texture, is routed through our handy Normal node. The ball has been rolled up and to the right. That puts out a mask that you can see in the Output viewer that is showing the Dot product output. Mixing that mask with the marbleized color and noise material/texture creates the final output for Material.001. As an exercise for the reader, giving the MatNode a color (instead of the grey) would create a specularity effect.

Using the Normal Node to create special Light-Sources



Using the Normal node as a special kind of lightsource.



With the â← Lightâ← from the Normal node.



Without the â← Lightâ← from the Normal node.

An obvious usage for the Normal node is to use it as a very special lightsource. In (*Using the Normal node as a special kind of lightsource*) the Orco coordinates are used, so the blue light rotates with the object rotation. If we would have used the Normal coordinates, we would have created a kind of object specific backlight. Global coordinates would give a more ordinary, stationary lamp.

The Fac value in the Mix node (the one with the caption â← Addâ←) determines the amount of added color.

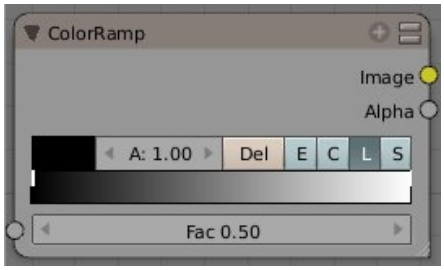
Material Convertor Nodes

As the name implies, these nodes convert the colors in the material in some way.

ColorRamp Node

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Convertors â ColorRamp



ColorRamp node
The ColorRamp Node is used for mapping values to colors with the use of a gradient. It works exactly the same way as a colorband for textures and materials, using the Factor value as a slider or index to the color ramp shown, and outputting a color value and an alpha value from the output sockets.

By default, the ColorRamp is added to the node map with two colors at opposite ends of the spectrum. A completely black black is on the left (Black as shown in the swatch with an Alpha value of 1.00) and a whitewash white is on the right. To select a color, LMB click on the thin vertical line/band within the colorband. The example picture shows the black color selected, as it is highlighted white. The settings for the color are shown above the colorband as (left to right): color swatch, Alpha setting, and interpolation type.

To change the hue of the selected color in the colorband, LMB click on the swatch, and use the popup color picker control to select a new color. Press â µ Enter to set that color.

To add colors, hold Ctrl down and CtrlLMB click inside the gradient. Edit colors by clicking on the rectangular color swatch, which pops up a color-editing dialog. Drag the gray slider to edit A:lpha values. Note that you can use textures for masks (or to simulate the old "Emit" functionality) by connecting the alpha output to the factor input of an RGB mixer.

To delete a color from the colorband, select it and press the Delete button.

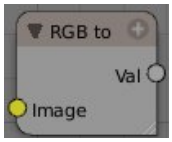
When using multiple colors, you can control how they transition from one to another through an interpolation mixer. Use the interpoloation buttons to control how the colors should band together: Ease, Cardinal, Linear, or Spline.

Use the A: button to define the Alpha value of the selected color for each color in the range.

RGB to BW Node

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Convertors â RGB to BW



RGB to BW node
This node converts a color image to black-and-white.

Connecting Output Socket
When you connect the output Val socket to an input socket that excepts an Image, Blender automatically inserts a ColorRamp node, to translate the output value to a material color.

Groups of Nodes

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Groups

Node maps can get quite complex. Blender allows you to group a set of nodes together to both save space and help you conceptualize what the net effect of a mini-map does to a material/image. This menu selection shows the names of the groups of nodes that you have defined. Select any one to add that group to the map.

Create a Node Group

A group is created by â § Shift-clicking all the nodes you want in the group, and then selecting Node -> Make Group (CtrlG). The group node is shown with a green bar, and the name is shown in an editable field (â § Shift-click on the name to enter EditMode and change the name to something that represents what that group does. The input sockets to the group is the input sockets of contained nodes, and akin for output(s).

Editing Node Groups

You can select the group and press â Tab (or select Node -> Edit Group) to enter/leave the group.

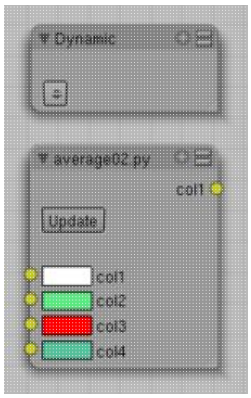
Ungroup Nodes

You can select the group and press AltG (or select Node -> Ungroup) to convert the grouped nodes back to normal.

Dynamic Nodes (PyNodes)

Panel: Node Editor → Material Nodes

Menu: â § ShiftA â Dynamic



PyNodes. Above one is in default state. One below has a script loaded into it.
Dynamic nodes allow you to write your own custom nodes. These nodes are written in Python. See API for more specific information.

To use a PyNode load the script of wanted node in Blender's text editor. After this you need to add Dynamic node to your node setup. The node added will contain a file selector. Use it to select the file loaded in text editor. After you have loaded it, it setups the node. If the code is valid, the node will be ready to use. If not, please check Blender's console to see what is wrong.

Note that if you make any changes to the script you need to press the Update button seen on the Dynamic node.

You can find PyNode recipes at the cookbook.

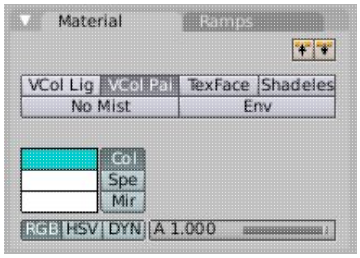
Blender features a built-in painter that allows you to paint your mesh all sorts of pretty colors. This section describes the Vertex Painter, which paints your base mesh by assigning colors to vertices, and blending those to give a face a color. Just like having your own virtual airbrush.

Setting Up

Mode: To be able to paint a mesh, you must select the object and go into Vertex Paint mode using the mode selector on the 3D Window header. Your cursor in the 3D window will switch to a cute little paintbrush.

Viewport Shading: To be able to **see** what you are painting, select Shaded as the Viewport Shading (also in the 3D window header). You don't have to be in Edit mode to paint, and although some find it distracting to see all those dots, it sometimes helps to see where the vertices actually are.

Textured Viewport Shading
If your 3D View is in Textured Viewport Shading mode, any Vertex Color paint layer will override a UV Texture layer, but will not affect rendering. In Shaded view, the correct view is shown: UV Textures override any Vertex Paint layer.
Detail: The more vertices you have, and the denser they are, the more detailed the paint effects will be. (The editing Subdivide button is really handy if you need it).

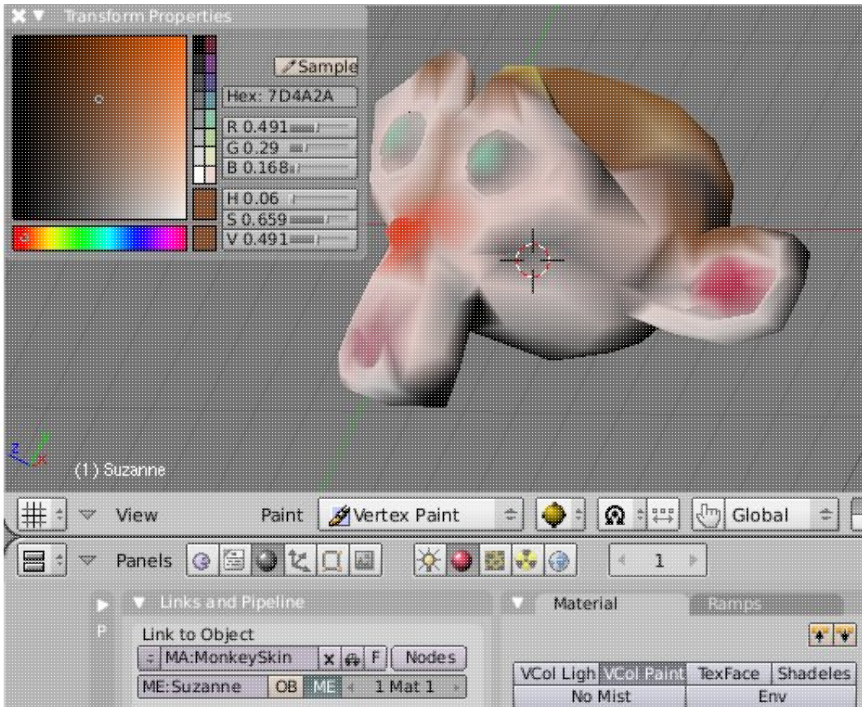


Enable Vertex Colors when Rendering
Render: You can paint all day inside Blender without having a material assigned to the object. However, you will not be able to Render without one. In the Buttons window, the Shading F5 buttons, assign a new material if the object does not have one. **Enable the VCol Paint button** in the material settings of the Shading buttons. This tells Blender to use the vertex colors you are going to paint (instead of the base material color) when it renders the image.

Multicolored lights: You can also tell Blender to use the vertex colors as a light source for the mesh when rendering; enable VCol Light and use the material Emit setting (on the Shaders panel) to vary the intensity of the light produced. Setting different colors for vertices, and then spinning the object, will make a warbling array of colors. If you want it to cast light and shadows on other objects, put a lamp inside of it and make it partially transparent (alpha less than 1.0).

Face Painting: By default, you will be painting the whole mesh. To only paint part of the mesh, press F for Faces when in Vertex Paint mode. The faces of the mesh will be outlined for you. Just like when editing a mesh, RMB clicking on a face, or â § ShiftRMB clicking, or doing a Border select with the LMB button selects only certain faces to paint. Doing a Border select with the RMB button excludes those faces from painting. Selecting some faces and pressing H Hides them from view and thus painting (but only for the current painting session; they are unhidden when you leave paint modeâ ←¦ or if you press AltH!).

Color Picker Applet



Press traNsform to have a floating color picker applet readily available. Use the color picker by RMB clicking on a predefined color, color bar and gradient, or LMB clicking the sliders, and/or eyedropper sampler. Choose a pretty color, and LMB click-drag over your mesh to apply the color. You may supply a

hissing sound when you click just to make it more entertaining.

Each time you paint a stroke without leaving the 3D window, Blender saves the stroke in the Undo buffer, so to undo a stroke, just press CtrlZ. Set the Opacity (density) and Size of your brush in the Paint panel (in a Editing (F9) context). Many other painting controls are available in that area; see the Reference Manual for details on the painting tool.

You may move your view about in the 3D window using normal 3D space controls. Note that the window shows you what the object will look like under current lighting conditions, so just like the real world, you may have to turn the object or add/move the lights to see what your paint job really looks like.

Vertex Paint Panel



In the Editing F9 buttons, in Vertex Paint mode, there will be a Paint panel that has many controls for your paintbrush. Use the RGB sliders to choose a color the same way as the floating applet. The panel also allows you to control the paint Opacity (how 'thick' the paint is applied), the Size of the brush, and how the paint is applied to existing colors. You can apply via Mix, Add, Subtract, Multiply, Filter, Lighter, or Darker. For example, if a vertex was purple, and you set your brush to Subtract Blue, painting that vertex would make it red, since purple minus the blue is red.

To make all vertices a consistent color, set the color and click Set Vert. By default, when you paint, the paint color spreads to all faces connected to the vertice and blends in based on the size of the face. Disable All Face and Vertex to uniformly paint a face. Enabling Normals shows you the incident light appearance; use this if you have Textures applied to the material that affect the Normal.

By default, holding down the LMB is just like holding down a spray can button; the more you hold, the more paint is applied. Disable Spray and each click sprays a little bit of paint but no more, and holding down the button has no effect.

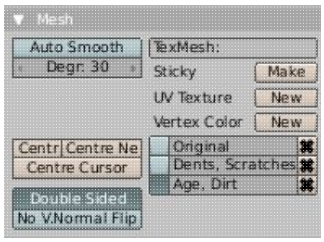
The bottom row of buttons allow you to uniformly multiply the color values, effectively increasing (a Multiply factor >1.0) or decreasing (factor <1.0). Change the value and click Set. Use Gamma correction for cross-platform image gamma correction.

Saving and Exporting

Whenever you save your .blend file, the vertex paint job is saved with it, inside the .blend file. There is nothing special you have to do.

You can also save your paint job as an external image (e.g. JPG or PNG) by baking the paint to a UV Image. You do this by unwrapping the painted mesh onto a flat surface, like carefully unwrapping a Christmas present and smoothing the paper out onto a tabletop. This process is called Baking the Vertex Paint and is discussed in UV Texturing.

Overpainting and Textures



Entering Vertex Paint mode creates a Vertex Color layer, named and indicated in the Editing buttons, Mesh Panel. You can create multiple layers of Vertex Paint by clicking the New button, located in the Editing buttons, Mesh Panel, to the right of the Vertex Colors label. Each layer is painted independently, and unpainted areas of an upper layer allow more original layers (shown higher in the list) to show through. Selecting a layer shows that painting in the 3D view and/or in Render output â← and deactivates (hides) all those that were overlaying it! Select a layer by clicking one or both of the buttons next to its name (left one for 3D view, right one for Render), and change the name by clicking on the name and entering something creative. The example has the Original paint job of the car, overlaid with some dents and scratches, which is further overlaid by the effects of aging and dirt (a light brown-gray dusting). When rendered, all three layers will be shown, because the third layer is currently selected.

You may also paint under a UV image by enabling the use of UV images (Material Texface button), assigning faces using the UV Face Select mode, and loading the image using the UV/Image editor. Any painting you do on one object's face that is UV-mapped does not affect the UV/image, but non-mapped faces will show the vertex colors you have painted. To make permanent mods to the UV image, use the painting tool in that window via Image->Texture Painting. Partially transparent UV images (with an alpha less than one) will allow the base vertex paint to show through.

Any texture (such as the cloud texture) that maps to the color of the material will also affect the vertex coloring.

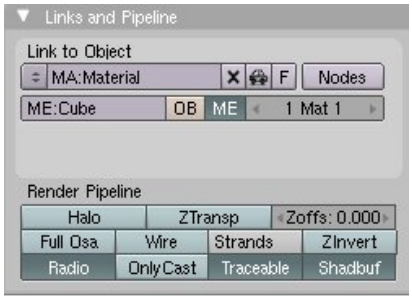
The end color of the material also depends on the amount of ambient light it receives and the color of that ambient light. If the material is partially transparent, then the color seen will also depend on the color of the objects behind it. The ulitmate color also depends on the color of lights (lamps, reflections, radiance/glow, and other vertex color lights) that shine upon it.

Blender provides a set of materials which do not obey the face-shader paradigm and which are applied on a per-vertex rather than on a per-face basis. These are called Halos because you can see them, but they do not have any substance. They are like little clouds of light; although they are not really lights because they do not cast light into the scene like a lamp.

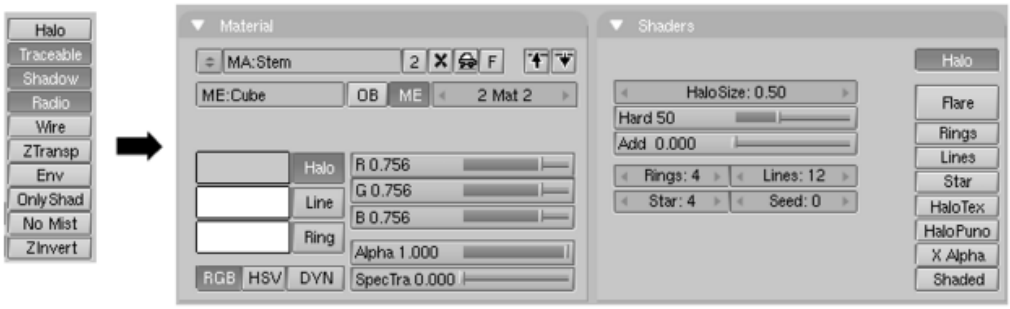
Halo Materials

Halos come in very handy when creating certain special effects, when making an object glow, or when creating a viewable light or fog/atmospherics around an actual light.

Options



Press F5 to display the Material buttons, and then press the Halo button on the Links and Pipeline Panel. The Material and Shaders panels change as shown below in *Halo buttons*.



Halo buttons

As you will see in the 3D View, the Mesh faces are no longer rendered. Instead just the vertex is rendered, since that is where each halo will originate. Halos can be hard to find in a crowded scene, so name it well for easy location in the outliner.

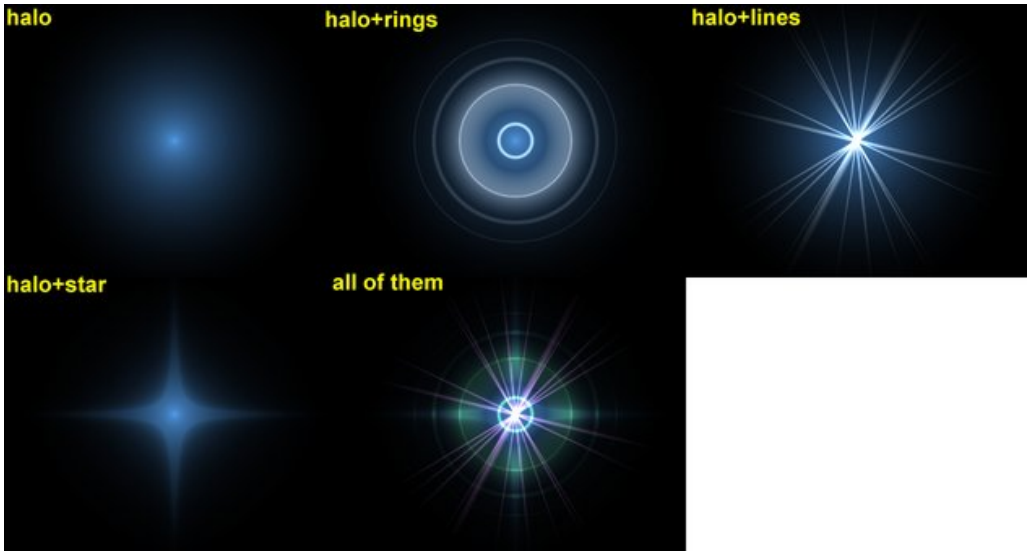
In the Material Panel, you now set three colors which, in standard material were Diffuse, Specular and Mirror colors are now relative to three different Halo characteristics:

Halo
: the color of the halo itself,

Line
: the color of any possible line you might want to add

Ring
: the color of any possible ring around the halo

You can not use color ramps. Lines, Rings and an assortment of special effects are available with the relevant toggle buttons, which include Flare, Rings, Lines, Star, Halotex, HaloPuno, X Alpha, and Shaded. *Halo Variations* shows the result of applying a halo material to a single vertex mesh.



Halo Variations
The halo size, hardness and alpha can be adjusted with the pertinent sliders. These are very similar to traditional material settings



The Add slider determine how much the halo colors are 'added to', rather than mixed with, the colors of the objects behind and together with other halos. By increasing Add, the Halo will appear to light up objects that move behind it or through the Halo field.

To set the number of rings, lines, and star points independently, once they are enabled with the relative Toggle Button, use the Num Buttons Rings:, Lines: and Star:. Rings and lines are randomly placed and oriented, to change their pattern you can change the Seed: Num Button which sets the random numbers generator seed.

Halo Texturing

By default, textures are applied to objects with Object coordinates and reflects on the halos by affecting their color, as a whole, on the basis of the color of the vertex originating the halo. To have the texture take effect *within* the halo, and hence to have it with varying colors or transparencies press the HaloTex button; this will map the whole texture to *every* halo. This technique proves very useful when you want to create a realistic rain effect using particle systems, or similar.

Another Option is Shaded. When shaded is enabled, the Halo will be affect by local light; a lamp will make it brighter and affect its diffuse color and intensity.

Examples

Dotmatrix display

Let's use a halo material to create a dotmatrix display.

- To begin, add a grid with the dimensions 32x16. Then add a camera and adjust your scene so that you have a nice view of the billboard.
- Use a 2D image program to create some red text on a black background, using a simple and bold font (if you are a lazy lizard [I hope this not offensive, I just like how it sounds!], you can just save the picture bellow on your hard driveâ←¦). *Dot matrix image texture.* shows an image 512 pixels wide by 64 pixels high, with some black space at both sides.



Dot matrix image texture.

- Add a material for the billboard, and set it to the type Halo. Set the HaloSize to 0.06 and when you render the scene you should see a grid of white spots.
- Add a Texture, then change to the Texture Buttons and make it an image texture. When you load your picture and render again you should see some red tinted dots in the grid.
- Return to the Material Buttons and adjust the sizeX parameter to about 0.5 then render again; the text should now be centered on the Billboard.
- To remove the white dots, adjust the material color to a dark red and render. You should now have only red dots, but the billboard is still too dark. To fix this enter EditMode for the board and copy all vertices using the â § ShiftD shortcut (take care not to move them!). Then adjust the brightness with the Add value in the MaterialButtons.



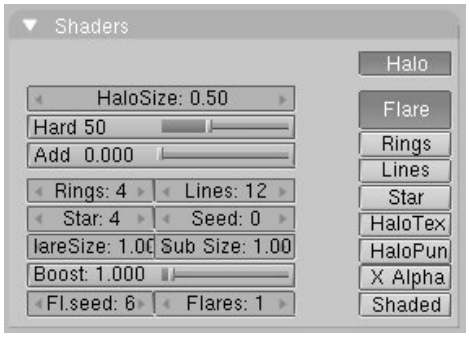
Dot Matrix display.

You can now animate the texture to move over the billboard, using the ofsX value in the Texture panel of the MaterialButtons. (You could use a higher resolution for the grid, but if you do you will have to adjust the size of the halos by shrinking them, or they will overlap. (*Dot Matrix display*).

Note about material indices
Halo materials only work when applied using the first material index. Any material(s) in a subsequent material index will not be rendered.

Lens Flares

Our eyes have been trained to believe that an image is real if it shows artifacts that result from the mechanical process of photography. *Motion blur*, *Depth of Field*, and *lens flares* are just three examples of these artifacts. The first two are discussed in the *chapter_rendering*; the latter can be produced with special halos. A simulated lens flare tells the viewer that the image was created with a camera, which makes the viewer think that it is authentic.



Lens Flare settings
We create lens flares in Blender from a mesh object using first the Halo button and then the Flare options in the Shaders Panel of the material settings. Try turning on Rings and Lines, but keep the colors for these settings fairly subtle. Play with the Flares: number and Fl.seed: settings until you arrive at something that is pleasing to the eye. You might need to play with FlareBoost: for a stronger effect (*Lens Flare settings*).

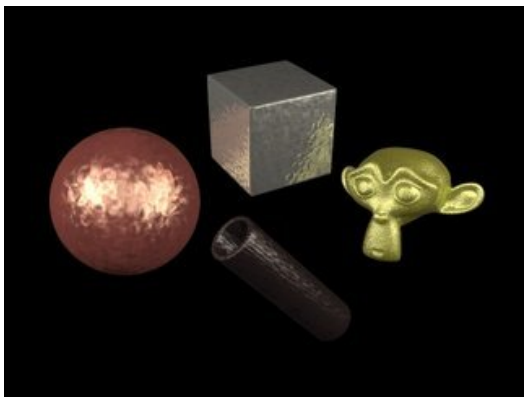
Note that this tool does not simulate the physics of photons travelling through a glass lens; it's just a eye candy.

Blender's lens flare looks nice in motion, and disappears when another object occludes the flare mesh. (see *Lens Flare* below).



Lens Flare

Introduction



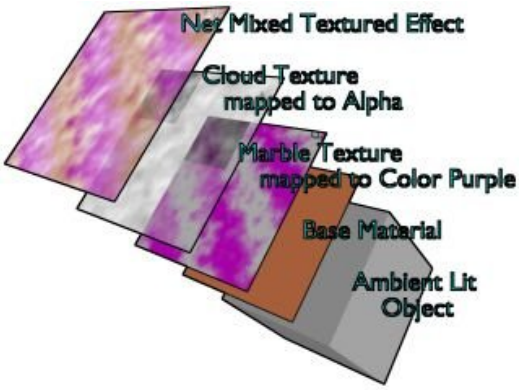
Some Metal Textures

The material settings that we've seen so far produce smooth, *uniform* objects, but such objects aren't particularly true to reality, where uniformity tends to be uncommon and out of place. In order to deal with this unrealistic uniformity, Blender allows the user to apply *textures* which can modify the reflectivity, specularity, roughness and other surface qualities of a material.

Textures fall into three primary categories:

- images,
- procedural textures (generated by a mathematical formula),
- environment maps (used to create the impression of reflections and refractions).

For an overview you may want to read our tutorial Using Textures.



Textures Layer on base Material

Textures are like additional layers on top of the base material. Textures affect one or more aspects of the object's net coloring. The net color you see is a sort of layering of effects, shown in this example image. The layers, if you will, are:

1. Your object is lit with **ambient** light based on your world settings.
2. Your base **material** colors the whole surface in a uniform color that reacts to light, giving different shades of the diffuse, specular, and mirror colors based on the way light passes through and into the surface of the object.
3. We have a primary **texture** layer that overlays a purple marble coloring.
4. We next have a **second cloud texture** that makes the surface transparent in a misty/foggy sort of way by affecting the Alpha value

5. These two textures are **mixed** with the base material to provide the net effect; a cube of purplish-brown fog.

This notion of using *more than one* texture, to achieve a combined effect, is one of the "hidden secrets" of creating realistic-looking objects. If you carefully "look at the light" while examining any real-life object, you will observe that the final appearance of that object is best described as the combination, in different ways and in different amounts, of several distinct underlying visual characteristics. These characteristics might be more (or less) strongly apparent at different angles, under different lighting conditions, and so forth. Blender allows you to achieve this in many ways. You can use "a stack of texture layers" as described in this section, or you can also use arbitrarily-complex networks ("noodles"...) of "texture nodes" as discussed here, the choice is yours.

Texture Buttons

Mode: All Modes

Panel: Shading/Textures Context

Hotkey: F6

Description

Once a new texture has been added to a material, it can be defined by switching to the Texture Buttons (F6) or  sub-context of the Shading context to obtain *Texture buttons*.

A new, empty texture Button Window presents two panels:

1. a Texture Preview and
2. a Texture panel, the latter with two tabs.

Texture Preview

Mode: All Modes

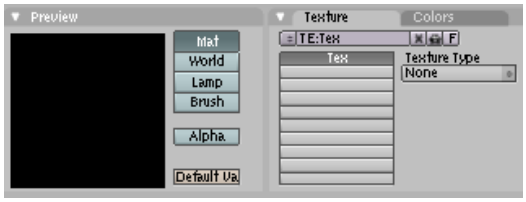
Panel: Shading/Textures Context → Preview

Hotkey: F6

Description

The texture preview panel provides a quick previsualisation of how the texture looks on its own, without mapping.

Options



Texture buttons
You can choose what kind of textures you are editing:

Material
: Edit the stack of textures linked to the active material

World
: Edit the stack of textures linked to the active world

Lamp
: Edit the stack of textures linked to the active lamp

Brush
: Edit the stack of textures linked to the active brush (in sculpt mode).

Alpha
: Show alpha in preview

Default Va
: Reset all the texture properties to their default values

Texture Channels

Mode: All Modes

Panel: Shading/Textures Context → Texture

Hotkey: F6

Description

This panel lets you manage the list of texture channels available for the particular shading data that you are working on (material, world or lamp).

Options

Standard datablock selector
: Choose, rename, unlink, automatically generate a name, or add a fake user for the active texture

Texture channel button
: Change the active texture channel



Texture Types.
Here you can also define the active texture's type. The available types of textures are:

Image
: Allows an image to be loaded and used as a texture. See Image Textures

EnvMap
: To simulate Reflections (and Refractions) without Raytracing. See Environment Maps.

Plugin
: Allows for loading an external piece of code to define the texture. See Texture Plugins.

Procedural
: The remaining options define 3D procedural textures, which are textures that are defined mathematically and are built into Blender. See Procedural Textures.

Texture Colors

Mode: All Modes

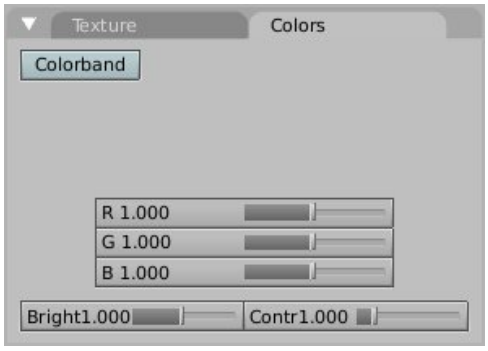
Panel: Shading/Textures Context → Colors

Hotkey: F6

Description

All textures may be modified by the Bright(ness) and Contr(ast) buttons in the Colors panel. All textures which posess RGB-Values - including Images and Environment Maps - may be modified with the RGB sliders. (*Texture Colors Panel*).

Options



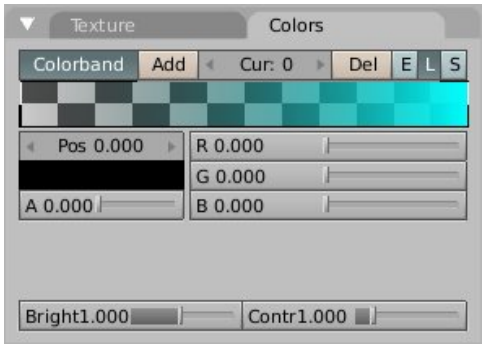
Texture Colors Panel.

R, G, B
: Tint the color of a texture by brightening each red, green and blue channel

Brightness
: Change the overall brightness/intensity of the texture

Contrast
: Change the contrast of the texture

Colorbands



Texture Colorband.

If intensity-only textures are used, the result is a black and white texture, which can be greatly enhanced by the use of colorbands. The colorband is an often-neglected tool in the Colors tab in the Texture Panel that gives you an impressive level of control over how procedural textures are rendered. Instead of simply rendering each texture as a linear progression from 0.0 to 1.0, you can use the colorband to create a gradient which progresses through as many variations of color and transparency (alpha) as you like (*Texture Colorband.*).

To use Colorbands, select a procedural texture, such as Wood. Click the Colorband button. The Colorband is Blender's gradient editor. Each point on the band can be placed at any location and can be assigned any color and transparency. Blender will interpolate the values from one point to the next.

For information on using the colorband UI controls, see Colorbands in the Ramps section of this manual.

Hints

- The alpha slider changes the transparency of the selected arrow of the colorband.
- If you use a colorband, the results of the texture are intensity and RGB. The alpha values deliver intensity, the colors RGB. Use the *NoRGB* button in the Materials context to calculate intensity from the RGB values.

See Also

- Ramps

Texture Channels

Mode: All Modes

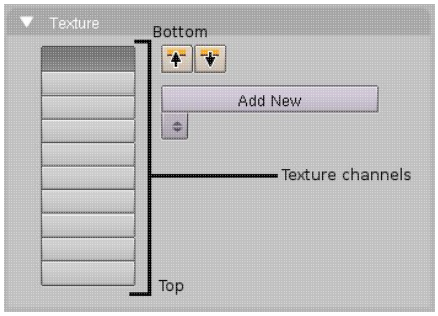
Panel: Shading/Material Context → Texture

Hotkey: F5

Description

A material can contain up to 10 Texture Channels, which are layered in a stack.
Each texture channel contains informations about the texture and its 'mapping' on the material, which describes how the texture is positioned on the object's geometry, and how it affects the material properties. A channel can be selected by clicking on it in the stack. By selecting it, you make it active.

Channels Options



Empty Texture Panel

Empty Channels

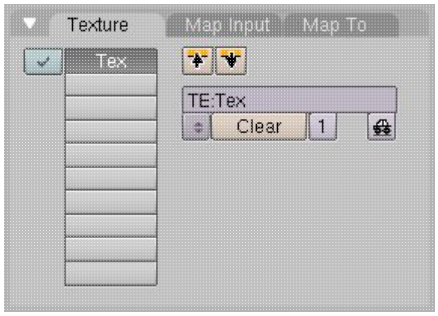
If there is no texture currently in the active texture channel, the following options are available:

Add New
: Add a new texture to the active texture channel

Select an existing texture

: Choose an existing texture from a list, to insert in the active texture channel.

Channels with Textures



Texture Panel with one texture
If the active texture channel contains a texture, further options are made available in the Texture panel, and in two additional panels, Map Input and Map To.

TE:
: The texture itself is designated by its name here, which you can edit here.

Select an existing texture
: Same as for #Empty Channels.

Clear
: Clears the texture from the active channel. The texture data still remains and can be added back to the same channel or another channel from the existing texture selection menu

Number of users (number field)
: The number of materials that use the texture. You can't create a "Single User" copy here, you have to do that in the Texture buttons (F6).

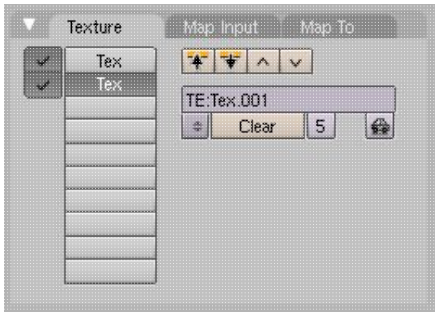
Auto (car icon)
: Automatically generates a (hopefully relevant) name for the texture

Texture Panel with one texture shows the texture with the two new panels. These panels are organized in the sequence in which the texture pipeline is performed:

- in the Map input panel, you choose how the active texture is mapped on the geometry
- in the Map to panel, you choose how the active texture affects the material properties

The channels stack

Each texture channel has its own individual texture mapping. Textures are layered upon each other and executed one after another from top to bottom. This means that for example the second channel overlays on top of the first channel.



Texture Panel with two textures

(De)activating a channel

By clicking the toggle at the left of the channel you enable or disable the correspondant texture in the stack.

Changing the order of the textures

By selecting a texture and clicking on the Up and Down arrows, you move the selected texture in the channels stack, so to change the overlay order.

Copying and pasting Texture Channels

By adding an existing texture to a channel, you create a link to that texture, but all the mapping options remain as they are. To copy all texture settings, including the contents of the Map Input and Map To panels, you can copy a given texture channel and paste it into another by using the copy and paste buttons.

Copying and Pasting is used when you want to combine two similar textures where each is slightly different from the other. Rather than manually duplicating the properties of one into another just copy and paste.



Copy (Arrow up)
: Copies the active texture channel's settings into a temporary buffer.

Paste (Arrow down)
: Pastes the copied texture channel information into the active texture channel.

Texture Map Input

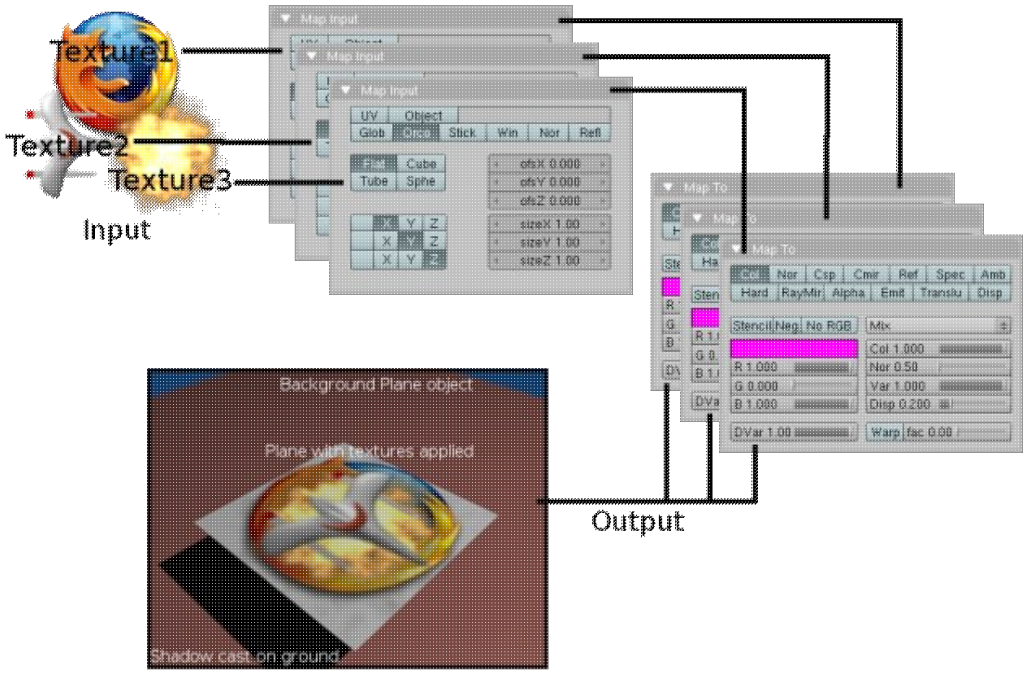
Mode: All Modes

Panel: Shading/Material Context → Map Input

Hotkey: F5

Description

Textures need mapping coordinates, to determine how they are applied to geometry. The mapping specifies how the texture will ultimately *wrap* itself to the object. For example, a 2D image texture could be configured to wrap itself around a cylindrical shaped object.



Conceptual texture pipeline

(*Conceptual texture pipeline*) is what is happening if you could "see" each panel(s) associated with each texture or image at the same time. However, only one panel is visible at any one time. To change a texture's Map Input properties you must first select the texture on the Texture panel.

Input Source

Mode: All Modes

Panel: Shading/Material Context → Map Input

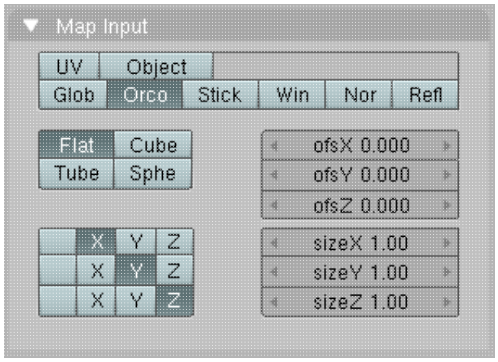
Hotkey: F5

Description

Mapping works by using a set of coordinates to guide the mapping process. These coordinates can come from anywhere, usually the object to which the texture is being applied to.

For UV or Object mapping, enter the name of the UV Texture or Object that you want to use as the orientation for the texture in the little blank field immediately to the right of the UV and Object buttons. The default name for a UV Texture is "UV Tex"; the exact name is found in the Materials panels where the UV Textures are listed. You can only specify one UV Texture for each mapped Texture image. If you want to layer UV Textures, you have to use multiple channels. See UV Unwrapping for more information on using UV Textures.

Options



Map Input Panel

UV

UV mapping is a very precise way of mapping a 2D texture to a 3D surface. Each vertex of a mesh has its own UV co-ordinates which can be unwrapped and laid flat like a skin. You can almost think of UV coordinates as a mapping that works on a 2D plane with its own local coordinate system to the plane on which it is operating on. This mapping is especially useful when using 2D images as textures, as seen in UV Mapping. You can use multiple textures with one set of UV Co-ordinates.

Some Modifiers Prevent UV Mapping
In fact, all modifiers affecting the *meshâ← s topology* (i.e. existing vertices, edges or faces) will prevent UV mapping, unless they are designed to handle this problem (for example, the Subsurf modifier manages nicely the UV maps, subdividing them as wellâ←¦). Others, like the Decimate modifier, affect the number of vertices, without knowing how to update accordingly the UV maps â← and thus completely prevent using UV mapping while they are in the modifiers stack.

To use an UV map to guide the texture placement, click *UV* and enter the name of the UV Texture in the *input field*. The name must match exactly and is case-sensitive. If the name does not match one of the existing UV Textures, the field will be red.

Then, for the Color of the texture, map the image texture to Col (if the image is the color/diffuse image). Map to Nor if the image is the bump map, etc. See Bump and Normal mapping for more info on those special kinds of images. See the next page for more info on Map To panel.

Object

Uses a child Object's texture space as source of coordinates. The Object name must be specified in the text button on the right. Often used with Empty objects, this is an easy way to place a small image as a logo or decal at a given point on the object (see the example below). This object can also be animated, to move a texture around or through a surface

Glob - Global

The scene's Global 3D coordinates. This is also usefull for animations; if you move the object, the texture moves across it. It can be useful for letting objects appear or disappear at a certain position in space.

Orco - Original Coordinates

"Original Co-ordinates" - The object's local texture space. This is the default option for mapping textures.

Stick

Uses a mesh's sticky coordinates, which are a form of per-vertex UV co-ordinates. If you have made *Sticky* co-ordinates first (F9 Mesh Panel, *Sticky* Button), the texture can be rendered in camera view (so called "Camera Mapping").

Win - Window

The rendered image window coordinates. This is well suited to blending two objects.

Nor - Normal
: Uses the direction of the surface's normal vector as coordinates. This is very useful when creating certain special effects that depend on viewing angle

Refl - Reflection
: Uses the direction of the reflection vector as coordinates. This is useful for adding reflection maps - you will need this input when Environment Mapping.

Hints

When choosing a mapping style consider the following questions:

- What is the source of the input Coordinates?
- Where is the top left corner?
- Where does the Texture begin?
- How large is the Texture and how often does it have to be repeated?
- Do we have multiple points to begin the texture, like at a cube?
- Will the Texture be rotated?

Mapping an image to the Window coordinates will cause the image to be tiled across the surface, without stretching, facing the camera without regard to the object's geometry.

Global coordinates, when mapped to Cube, tile the virtual space with the image. Wherever the object happens to be, it gets the texture that exists in that space.

Textures scale along with the object. If you scale an object larger, the image or procedural texture is scaled up as well. To make more patterns of the texture repeat across the surface, either increase the image repeat value (if an image) or increase the SizeXYZ values in the Map Input panel. This can be a pain if you have a texture such as a brick, and you make the wall longer; the bricks will stretch out. You can either use a procedural plug-in texture that figures out the repeats for you, or get out the calculator and divide the "true" texture size by the actual wall size to determine how many repeats there should be.

Use multiple texture channels of the same texture, each Sized differently, to create the natural patterns found in nature (principle of the fractals).

To make them completely independent from the object's scale, either use Global mapping or Object, referencing an empty at the object's center with a scale of 1 (preferably parented to the object if you are animating it.).

2D to 3D Mapping

Mode: All Modes

Panel: Shading/Material Context → Map Input

Hotkey: F5

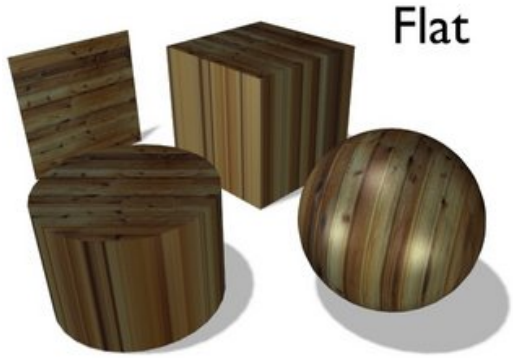
Description

The Image texture is the only true 2D texture, and is the most frequently used and most advanced of Blender's textures. Because images are two-dimensional, the way in which the 2D texture coordinate is translated to 3D must be specified in the *mapping* buttons.

Options

Depending on the overall shape of the object, one of these types may be more useful than others.

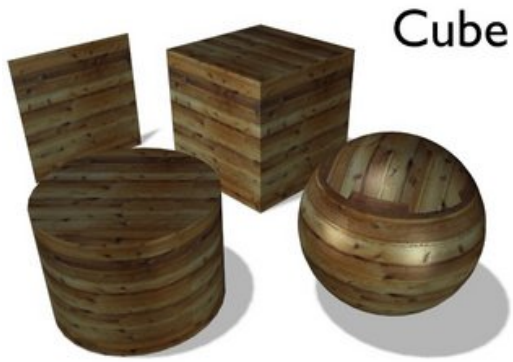
If the UV button (see #Input_Source; *Others also?*) is enabled you can imagine all these buttons except for the Flat button to be disabled.



Flat Mapping.

Flat

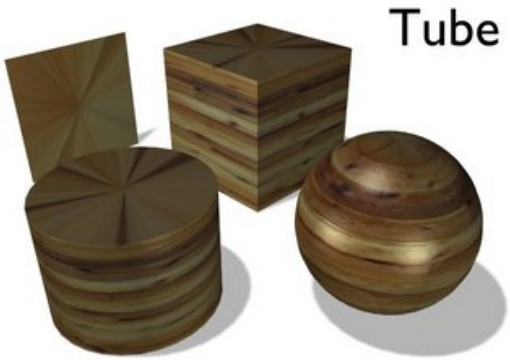
Flat mapping gives the best results on single planar faces. It does produce interesting effects on the sphere, but compared to a sphere-mapped sphere the result looks flat. On faces that are not in the mapping plane the last pixel of the texture is extended, which produces stripes on the cube and cylinder.



Cube Mapping.

Cube

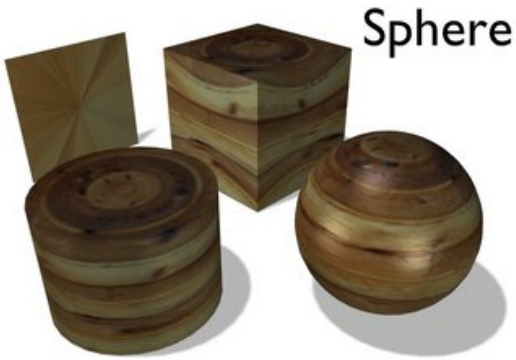
Cube mapping often gives the most useful results when the objects are not too curvy and organic (notice the seams on the sphere).



Tube Mapping.

Tube

Tube mapping maps the texture around an object like a label on a bottle. The texture is therefore more stretched on the cylinder. This mapping is of course very good for making the label on a bottle or assigning stickers to rounded objects. However, this is not a cylindrical mapping so the ends of the cylinder are undefined.



Sphere Mapping.

Sphere

Sphere mapping is the best type for mapping a sphere, and it is perfect for making planets and similar objects. It is often very useful for creating organic objects. It also produces interesting effects on a cylinder.

Coordinate Offset, Scaling and Transformation

Mode: All Modes

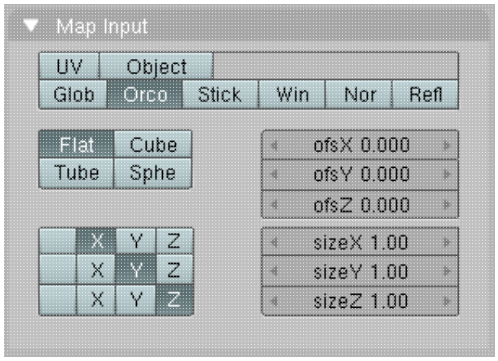
Panel: Shading/Material Context → Map Input

Hotkey: F5

Description

For extra control, the texture space can also be tweaked, to move, scale and flip the apparent texture on its axes

Options



Map Input panel.

ofsX, ofsY, ofsZ â← Offset
: The texture co-ordinates can be translated by an offset. Enlarging of the Ofs moves the texture towards the top left.

sizeX, sizeY, sizeZ â← Size
: The scale of the texture space. Enlarging the texture space will make the apparent texture scale smaller. The texture is as often repeated (if set up as repeating) as sized here. You can also use these settings to stretch the texture, by entering different values in each numeric fields. And finally, you can mirror a texture, by setting a negative value as scale of the axis along which you want to mirror.

[], X, Y, Z (x3) â← Axes
: These buttons allow you to change the mapping of axes between the textureâ← s own coordinate system, and the mapping system you choose (Orco, UV, etc.).
More precisely, to each axis of the texture (from top to bottom, X, Y and Z) corresponds one row of four buttons, that allow you to select to which axis in the mapping system it maps â← the first void button mapping to nothing! This implies several points:
 - ◊ For 2D textures (imagesâ←¦), only the first two rows are relevant, as they have no Z data!
 - ◊ You can rotate a quarter turn a 2D picture by setting the first row (i.e. X texture axis) to Y, and the second row (Y texture axis) to X.
 - ◊ When you map no texture axis (i.e. the three â← voidâ← buttons are set), youâ← ll get a solid uniform texture, as you use *zero dimension* (i.e. a dot, or pixel) of it (and then Blender extends or repeats this pointâ← s color along all axes!).
 - ◊ When you only map one texture axis (i.e. two â← voidâ← buttons are enabled), youâ← ll get a â← stripedâ← texture, as you only use *one dimension* (i.e. a line of pixel) of it, and then Blender stretches this line along the two other axesâ←¦
 - ◊ The same goes, for 3D textures (i.e. procedural ones), when one axis is mapped to nothing, Blender extends the plan (â← sliceâ←) along the relevant third axisâ←¦

 So, all this is a bit hard to understand and master. Fortunately, you do not have to change these settings often, except for some special effectsâ←¦ Anyway, the only way to get used to them is to practice!

3D View texture space transform

Mode: Object Mode

Hotkey: T

Description

The texture space can also be transformed interactively in the 3D View, just like moving or scaling an object. This determines the area that the texture uses to define its co-ordinates.

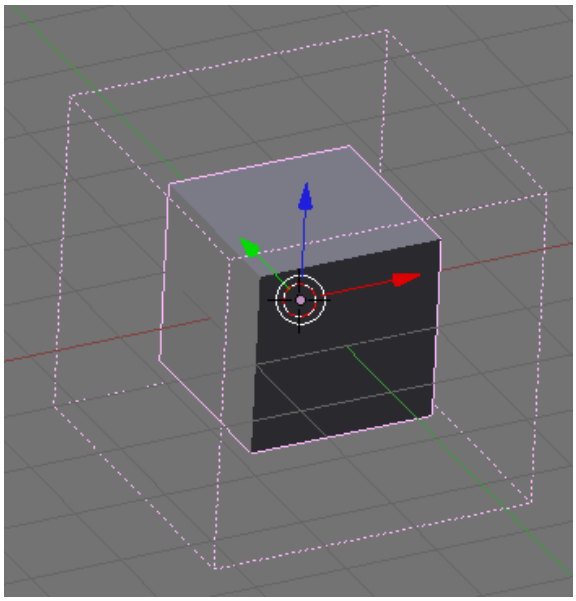
Options



Texture Space popup menu
Press T while an object is selected to get the Texture Space popup menu. The following options are available:

Grab/Move
: Moves the objects' texture space

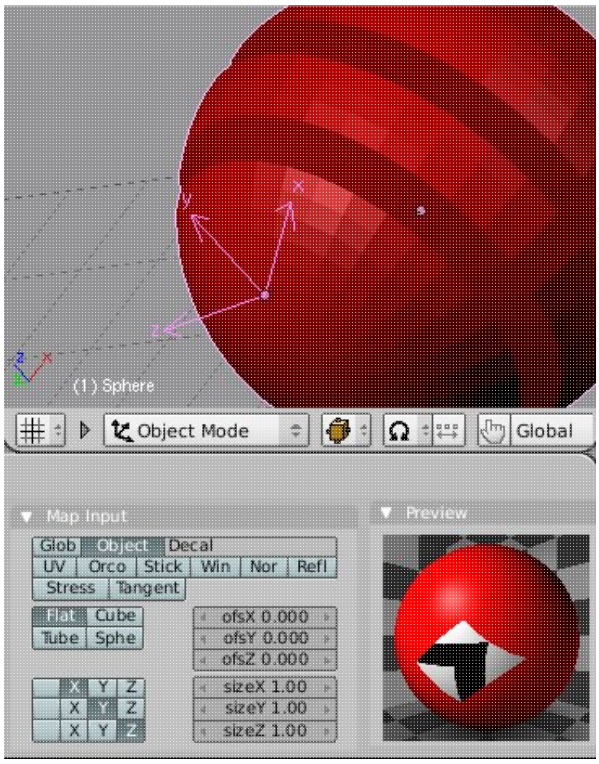
Scale
: Scales the object's texture space



Scaled Texture Space box.
Reset Texture Space
This manual editing overrides the automatic calculation of texture space based on the object's local co-ordinates and size. To return to automatic texture space calculation, enable the AutoTexSpace button in the Link and Materials Panel in the Editing context (F9)

Examples

Position a Decal on a Mesh



Another common use of Map Input, briefly mentioned above, is placing a decal somewhere on a mesh:

- Insert an Empty into your scene, positioned close to the surface of your mesh where you want the center of the decal to be, and oriented with its X-Y axis as you want your decal to appear (the Z axis should point away from the mesh, i.e. normal to the mesh).
- Assign an image texture to the mesh, loading the decal image, and set UseAlpha and ClipCube in the Map Image panel.
- In the Map Input field, select Object and in the blank field to the right of the button, enter the name of the Empty, which by default is "Empty"

In the example to the right, an Empty named "Decal" is just above the surface of the red ball, with the Z pointing away from the ball. In the material settings for the ball, an image of an arrow is mapped to the location and XY orientation of the empty.

You may Scale the empty to make the decal larger or smaller, and/or use the Size sliders in the Map Input panel. You can tweak the position of the decal by moving the empty or using the Offset sliders.

Texture Map To

Mode: All Modes

Panel: Shading/Material Context → Map To

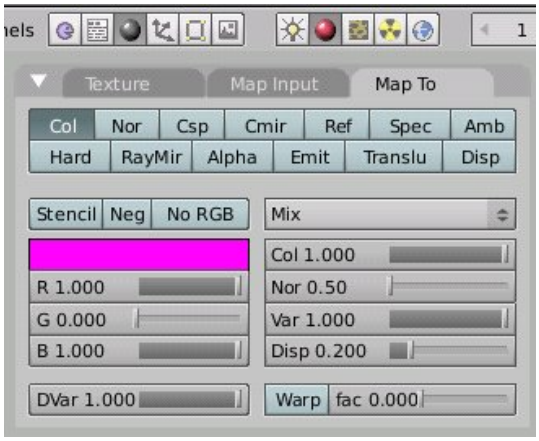
Hotkey: F5

Description

Not only can textures affect the color of a material, they can also affect many of the other properties of a material. The different aspects of a material that a texture influences are controlled in the Map To panel.

The result is also dependent from the type of input value. A texture may carry Intensity Information (0-255), Transparency (0.0-1.0), RGB Color (3 channels of 0.0-1.0) or Normal Vectors (either Bump or real Normals, see Section Bump and Normal Maps).

Options



Map To Panel.

Aspect

Some of the following options use three-state toggle buttons, meaning that the texture can be applied as positive or negative to some aspect of the material, or not at all. All of these buttons are independent.

Col (on/off)
: Influences the Material's RGB color

Nor (+/-/off)
: Commonly called bump mapping, this alters the direction of the surface normal. This is used to fake surface imperfections or unevenness via bump mapping, or to create reliefs.

Csp (on/off)
: Influences the Specular color, the color of the "reflects" created by the lamps on a glossy material.

Cmir (on/off)
: Influences the mirror color. This works with environment maps and raytraced reflection

Ref (+/-/off)
: Influences the amount of diffuse reflection.

Spec (+/-/off)

Influences the amount of specular reflection.

Amb (+/-/off)
: Influences the amount of Ambient light the material receives

Hard (+/-/off)
: Influences the specular hardness amount. A DVar of 1 is equivalent to a Hardness of 130, a DVar of 0.5 is equivalent to a Hardness of 65.

RayMir (+/-/off)
: Influences the strength of raytraced mirror reflection

Alpha (+/-/off)
: Influences the Opacity of the material. See Use Alpha for Object Transparency. Also use ZTransp for light and if combining multiple channels.

Emit (+/-/off)
: Influences the amount of light Emitted by the material

Translu (+/-/off)
: Influences the Translucency amount

Disp (+/-/off)
: Influences the Displacement of vertices, for using Displacement Maps.

Channel Filter

Stencil
: The active texture is used as a mask for all following textures. This is usefull for semitransparent textures and "Dirt Maps". See the example below (Stencil). Black sets the pixel to "untexturable".

Neg
: The effect of the Texture is negated. Normally white means on, black means off, Neg reverses that.

No RGB
: With this option, an RGB texture (affects color) is used as an Intensity texture (affects a value)

Color Swatch
: If the texture is mapped to Col, what color is blended in according to the intensity of the texture? Click on the swatch or set the RGB sliders

DVar
: Destination Value (not for RGB). The value with which the Intensity texture blends with the current value. Two examples:

- The Emit value is normally 0. With a texture mapped to Emit you will get maximal effect, because DVar is 1 by default. If you set DVar to 0 no texture will have any effect.
- If you want transparent material, and use a texture mapped to Alpha, nothing happens with the default settings, because the Alpha value in the Material panel is 1. So you have to set DVar to 0 to get transparent material (and of course ZTransp also). This is a common problem for beginners. Or do it the other way round - set Alpha to 0 and leave Dvar on 1. Of course the texture is used inverted then.

Blending Mode

How this channel interacts with other channels below it. See the Compositing Mix node for information and examples on the effect of each mixing mode.

Mix
Add, Subtract, Multiply, Divide
Screen, Overlay, Difference
Darken/Lighten
: Hue, Saturation, Value

Impact Sliders

Col
: The extent to which the texture affects colour

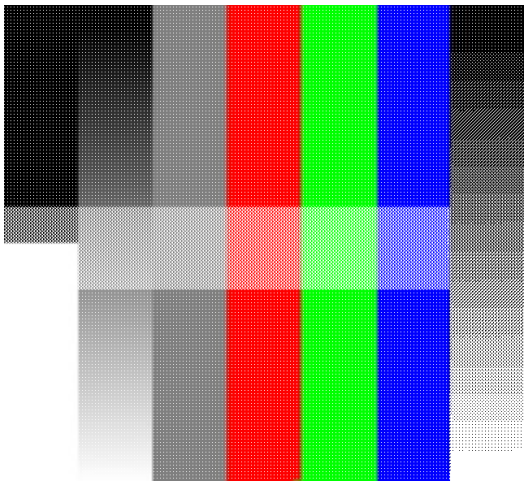
Nor
: The extent to which the texture affects the normal. Affects Normal, Bump and Displacement Maps.

Var
: The extent to which the texture affects the other values

Disp
: The extent an intensity texture changes the displacement. See section Displacement Maps.

Warp and Fac (factor)
: Distort subsequent textures to give an illusion of shape. See (Warp).

Examples





The input file. Vertical: Black, White, Black/White blend, 50% Grey, RGB, Black/Alpha blend. The brighter horizontal bar has a transparency of 50%.
The *input file* demonstrates the impact of different input values. The color of the underlying plane is magenta (R=1.0, G=0.0, B=1.0). The Map To → Col color of the texture is yellow (R=1.0, G=1.0, B=0.0).



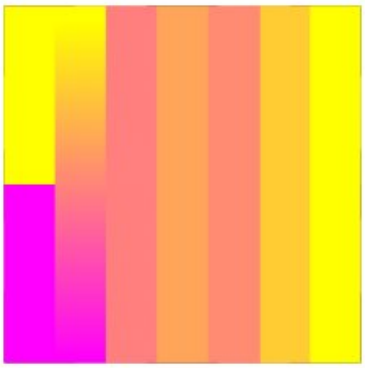
Use Alpha
Normally the color of the texture is opaque. If you Use Alpha with the image texture the alpha information of the image is evaluated. This will not make the material transparent, only the texture! So these pixels show the underlying material color (*Use Alpha*).

No RGB shows the result of the NO RGB option. The RGB texture is used as an Intensity texture. White produces the "Map To" color.

Neg invertes the respective values.



No RGB



No RGB and Neg

Hints

Every texture is opaque to the textures above it. This is no problem, if one texture is for example mapped to color, the other mapped to alpha etc. If you want to give a Mesh multiple materials see instead Multiple Materials.

Stencil

Mode: All Modes

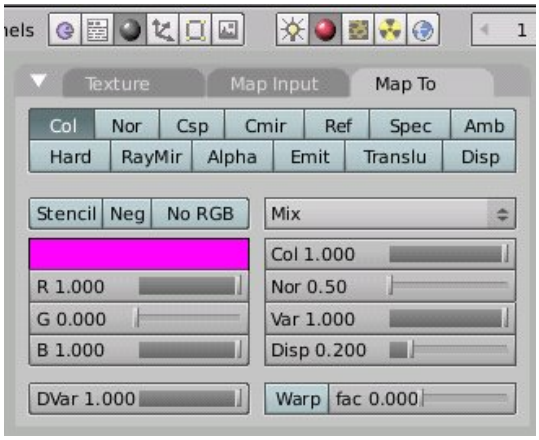
Panel: Shading/Material Context → Map To

Hotkey: F5

Description

The Stencil mode works similar to a layer mask in a 2D program. The effect of a stencil texture can not be overridden, only extended. You need an intensity map as input.

Options



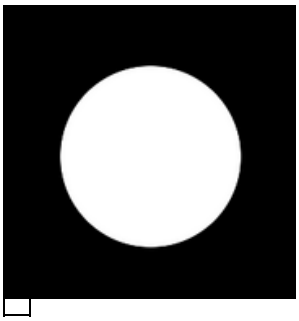
Map To Panel.

Stencil
: The active texture is used as a mask for all following textures. Black sets the pixel to "untexturable".

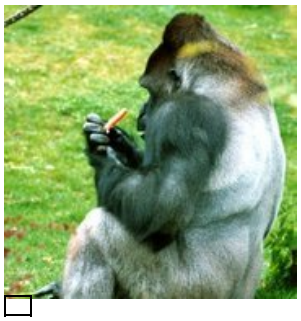
Examples

The Stencil mode works similar to a layer mask in a 2D program. The effect of a stencil texture can not be overridden, only extended. You need an intensity map as input.

Where the mask is black the following textures have no effect. Stencil needs intensity as input, so you have to use No RGB if you would like to use either images without alpha or a texture with a colorband (e.g. Blend Textures with a colorband).



The mask



The texture



The result



Stencil map with a radial blend.
You can blend two textures if you use a smooth blend texture as stencil map (*Stencil map with a radial blend*).

Warp

Mode: All Modes

Panel: Shading/Material Context → Map To

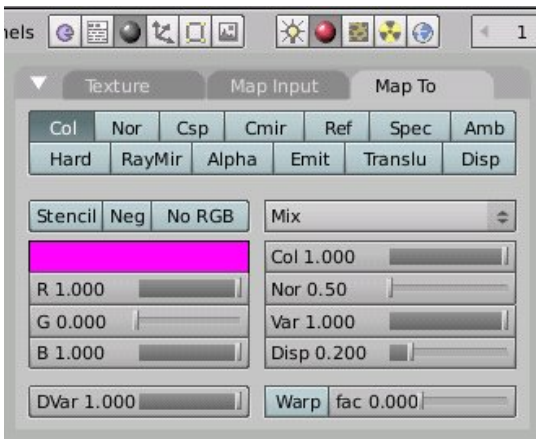
Hotkey: F5

Description

The option Warp allows textures to influence/distort the texture coordinates of a next texture channel.

The distortion remains active over all subsequent channels, until a new Warp has been set. Setting the fac at zero cancels out the effect.

Options



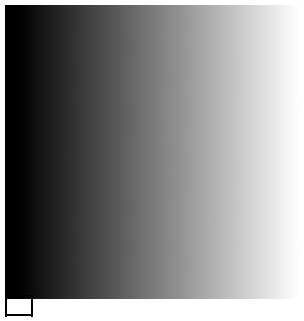
Map To Panel.

Warp
: Enable and disable the warp distortion

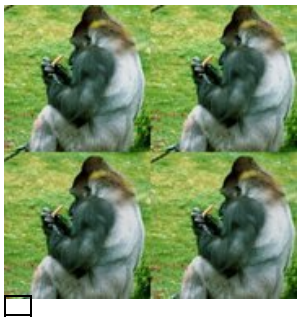
Fac
: The amount of distortion

Examples

The following example warps a gorilla texture based on a simple blend texture:



The Blend Texture



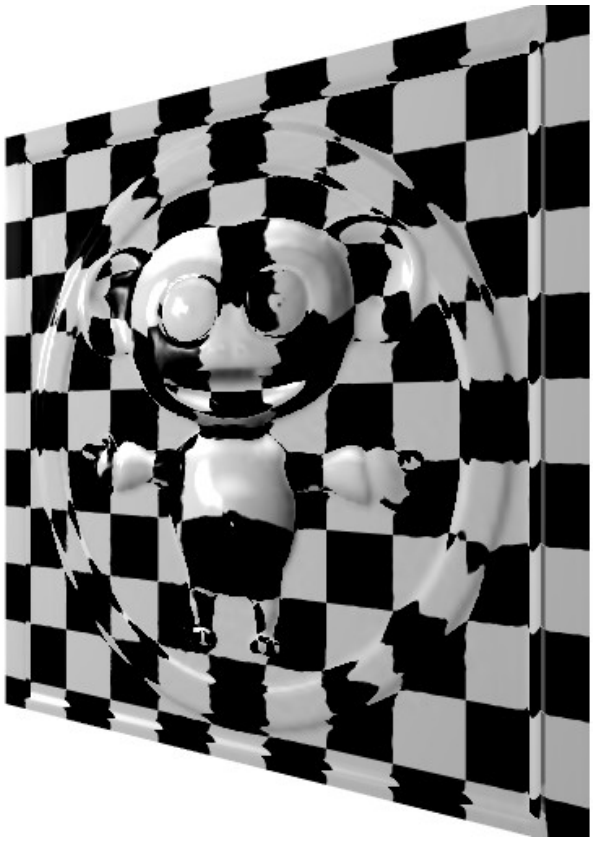
The Texture to warp



The warped Result



The same Texture for Normal and Warp mapping.
In this example, the normal map of Cornelius (*The same Texture for Normal and Warp mapping*) was used as normal map as well as warp texture in channel 1. The checkerboard texture is used in channel 2.



Cornelius as a Warp Factor

Bump and Normal Maps

Mode: All Modes

Panel: Shading/Texture Context → Image

Hotkey: F6

Description

Normal Maps and Bump Maps both serve the same purpose: they simulate the impression of a detailed 3D surface, by modifying the shading as if the surface had lots of small angles, rather than being completely flat. Because it's just modifying the shading of each pixel, this will not cast any shadows and will not obstruct other objects. If the camera angle is too flat to the surface, you will notice that the surface is not really shaped.

Both bump maps and normal maps work by modifying the normal angle (the direction pointing perpendicular from a face), which influences how a pixel is shaded. Although the terms normal map and bump map are often used synonymously, there are certain differences.

Bump maps
: These are textures that store an **intensity**, the relative height of pixels from the viewpoint of the camera. The pixels seem to be moved by the required distance in the direction of the face normals. You may either use greyscale pictures or the intensity values of a RGB-Texture (including images).

Normal maps
: These are images that store a **direction**, the direction of normals directly in the RGB values of an image. They are much more accurate, as rather than only simulating the pixel being away from the face along a line, they can simulate that pixel being moved at any direction, in an arbitrary way. The drawbacks to normal maps are that unlike bump maps, which can easily be painted by hand, normal maps usually have to be generated in some way, often from higher resolution geometry than the geometry you're applying the map to.

 Normal maps in Blender store a normal as follows:

 ◊ Red maps from (0-255) to X (-1.0 - 1.0)
 ◊ Green maps from (0-255) to Y (-1.0 - 1.0)
 ◊ Blue maps from (0-255) to Z (0.0 - 1.0)

 Since normals all point towards a viewer, negative Z-values are not stored (they would be invisible anyway). In Blender we store a full blue range, although some other implementations also map blue colors (128-255) to (0.0 - 1.0). The latter convention is used in "Doom 3" for example.

Workflow

The steps involved in making and using Bump and Normal Maps is:

1. Model a highly detailed ("hi-poly") model
2. Bake the Bump and/or Normal maps
3. Make a low-poly, less detailed model
4. Map the map to the low-poly model using a common coordinate system

Consult the Modeling section for how to model a highly detailed model using the Mesh tools. How much detail you put in is totally up to you. The more ridges and details (knobs, creases, protrusions) you put in, the more detailed your map will be.

Baking a map, simply put, is to take the detail of a high polygon mesh, and apply it to a similar object. The similar object is identical to the high-poly mesh except with less vertices. Use the Render Bake feature in Blender to accomplish this.

Modeling a low-poly using Blender's Mesh editing tools. In general, the same or similar faces should exist that reflect the model. For example, a highly detailed ear may have 1000 faces in the high-poly model. In the low-poly model, this may be replaced with a single plane, oriented in the same direction as the detailed ear mesh. *(Tip:* Blender's multi-resolution mesh modeling feature can be used to good effect here.)

Mapping is the process of applying a texture to the low-poly mesh. Consult the Textures Map To section for more information on applying a texture to a mesh's material. Special considerations for Bump and Normal Maps is:

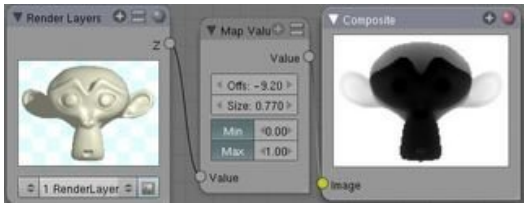
- When using a Bump map, map the texture to *Nor* and enable *No RGB.*
- When using a Normal map, map the texture to *Nor*

The coordinate systems of the two objects must match. For example, if you bake using a UV map of the high-poly model, you must UV map the low poly model and line up its UV coordinates to match the outline of the high-poly image (see UV unwrapping to line up with the high-poly map edges.

Examples

To set up the scene, position an orthographic camera about 10 units from a monkey. UV unwrap the Monkey; in this example we did an easy *Project from View (Bounds)*. Scale the camera so that the monkey fills the frame. Set your format resolution to 512x512 (square textures render faster in game engines, usually). No lights or anything else is required.

Use Render Bake to create a Normal map.

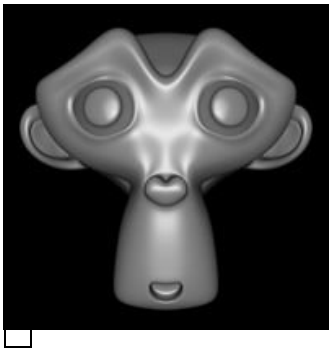


Composite noodle to make Bump Maps

Then, enable *Do Composite* and set up the compositing noodle shown to the right. This is your Bump Map. You make a bump map by using that same orthographic camera in the composite noodle shown. The front of the model is 9.2 units from the camera, and the visible parts are 1.3 units deep (1/1.3=0.77). If you save your image as a JPG, you will have 24-bit depth. If you save it as a half OpenEXR image, 16-bit depth; as a full OpenEXR image, 32-bit depth. Alternatively, you can use the ZUtilz Plugin, however, it only gives you an 8-bit depth and the resolution of the Bump Map may be too small for the relatively great range of Z-Values.

Now we are going to look at some examples. First the render of "Suzanne" (*Suzanne Render*). The second picture shows Suzanne's Normal Map (made with Blender's Normal Baking system. The rightmost picture in the top row shows the *Bump Map of Suzanne*.

Both maps here are used as textures on a plane, first the Normal Map (*Render of the Normal Map*), then the Bump Map (*Render of the Bump Map*). In both cases the camera stayed in the same position in which the Maps were made (perpendicular to the plane).



Suzanne Render



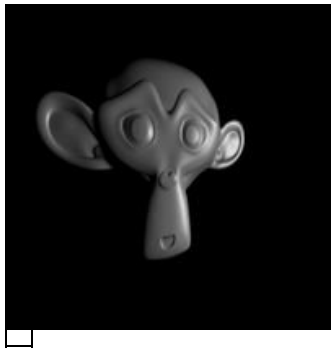
Normal Map of Suzanne



Bump Map of Suzanne



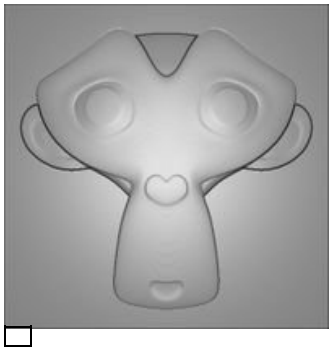
Side view render of the Normal Map applied to a plane



Side view render of a Normal Map made with an Ortho camera, applied to a plane. The same camera position as in the previous picture, less perspective distortion.



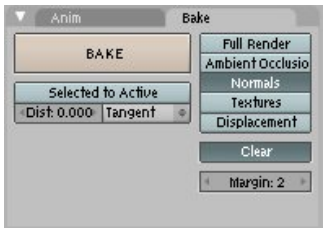
Render of the Normal Map applied to a plane, perpendicular to the surface.



Render of the Bump Map applied to a plane.
The *Render of the Normal Map* is only pseudo 3D. You can't look at the side of the head (*Side view render*).

If you use an Ortho camera to create the Normal Map, you will get less perspective distortion (*Render with an Ortho camera*).

Baking Options



The Render Bake panel
Render baking is hyperlinked here. In summary, there Blender supports normal map baking, amongst other bakeable things:

- Tangent: Bakes a normal map that is independent of view, and is thus the best choice for animation.
- Object: Bakes normals in the object's coordinates, and can be moved, but not deformed.
- Camera: Bakes normals with the old method, meaning you cannot deform or move the object.
- World: Bakes normals with coordinates, but object can't be moved or deformed.

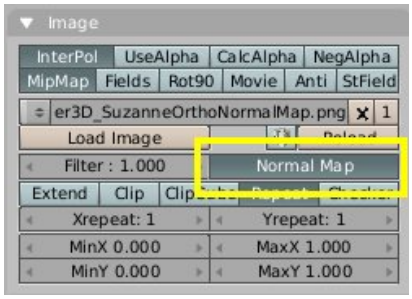
Other options include:

- Selected to active: This option will bake the selected mesh's normals

to the active mesh (e.g. the high poly mesh is the selected and is baked to the low-poly mesh which is the active).

- Distance: A parameter that controls the max distance from one object to another.
- Clear: Wipes the current image clear of any data.
- Margin: The amount (in pixels) of which to extend the normal map.

Using And Creating Normal Maps



Normal Map Button in the Image Panel
Normal and Bump Maps are simple to use. Make sure you apply the Texture in the Material buttons Map To panel to Nor. The strength of the effect is controlled with the NumButton Nor on the same Panel.

If you want to use a Normal Map, you have to set the Normal Map button in the Image Panel in the Texture buttons F6 (*Normal Map Button*).

Since only the normals get affected during render, you will not get shadow or AO or other '3D' effects. It's still just a texture.

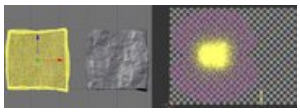
To bake normal maps is a different process, but simple.

1) Make a mesh.

2) Create a low polygon version of the object either by lowering the multires level (if you are using multires), by using a decimate modifier, or by just modeling the same mesh with less vertices. For simplicity, we'll say the low polygon object= mesh L and the high polygon object= mesh H.

3)To make this all work, mesh L needs to have a UV image. Assign seams and unwrap to a new image of the desired size (with UV test grid enabled. This will be explained in the "Common Problems" section further down the page).

You'll notice a new "32 bit float" option under the new image panel. This

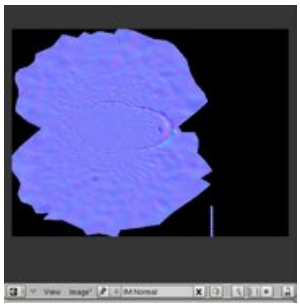


The low-poly mesh unwrapped
option creates your image in 32 bits, which is useful when generating displacement maps. If saving the 32-bit image, it must be saved to a format supporting 32 bits, such as OpenEXR.

4)Before you bake the normal maps, both meshes must be in the exact same location, so mesh L must be moved back to mesh H's location with ALT + G.

5)If the "Selected to active" button is enabled (which it should), you will need to select first mesh H, and then mesh L.

6) Go to the "Scene" tab (F10) and then to the bake tab (next to "anim"). Select "Normals", and customize your options according to your needs. Then press the big shiny "Bake" button and watch your map appear in all it's glory.



The normal map after baking

7) Nope, you're not done yet. We still need to add the normal map as a texture. The next step is to go into the materials tab (F5) and apply a new texture. Select the type as "Image",

and select your normal map texture (hopefully you named it) with the arrows.

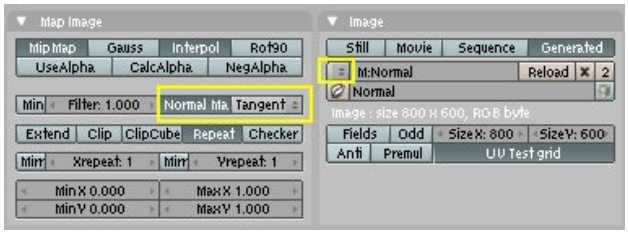


Image selected
8) Back in the main material panel under "Map Input", change the texture coordinates to UV. Under "Map To", check "Nor" and deselect "Col". Set the "Nor" value to 1 (Or a different value, depending on how strong you want it). You're done!

Displacement Maps

Panel: Material buttons → Map To

Hotkey: F5

Description

Displacement mapping allows a texture input to manipulate the position of vertices on rendered geometry. Unlike Normal or Bump mapping, where the shading is distorted to give an illusion of a bump (discussed on the previous page), Displacement Maps create real bumps, creases, ridges, etc in the actual mesh. Thus, the mesh deformations can cast shadows, occlude other objects, and do everything that changes in real geometry can do.

Options



Settings for a Displacement Map.
The strength of the displacement is controlled with the number fields Disp and Nor.

- If a texture provides only normal information (e.g. Stucci), vertices move according to the texture's normal data. The normal displacement is controlled by the Nor slider.
- If a texture provides only intensity information (e.g. Magic, derived from color), vertices move along the directions of their normals (a vertex has no normal itself, it's the resulting vector of the adjacent faces). White pixels move outward in the direction of the normal, black pixels move in the opposite direction. The amount of displacement is controlled with the Disp slider.

The two modes are not exclusive. Many texture types provide both information (Cloud, Wood, Marble, Image). The amount of each type can be mixed using the respective sliders. Intensity displacement gives a smoother, more continuous surface, since the vertices are displaced only outward. Normal displacement gives a more aggregated surface, since the vertices are displaced in multiple directions.

The depth of the displacement is scaled with an object's scale, but not with the relative size of the data. This means if you double the size of an object in object mode, the depth of the displacement is also doubled, so the relative displacement appears the same. If you scale inside editmode, the displacement depth is not changed, and thus the relative depth appears smaller.

Hints

Displacement maps move the render faces, not the physical mesh faces. So, in 3D View the surface may appear smooth, but render bumpy. To give a detailed surface, there has to be faces to displace and have to be

very small. This creates the tradeoff between using memory and CPU time versus render quality.

From best to worst, displacement works with these object types using the methods listed to control the render face size:

- **SubSurfaced Meshes**

 Render face size is controlled with render subsurf level. Displacement really likes smooth normals.

- **Manually (editmode) subdivided meshes**

 Control render faces with number of subdivides. (This can be combined with the above methods.) Displaces exactly the same Simple Subsurf, but slows editing down because of the OpenGL overhead of drawing the extra faces. (You can't turn the edit subdivide level down this way).

- **Meta Objects**

 Control render faces with render wiresize. Small wire == more faces.

The following are available, but currently don't work well. It is recomended that you convert these to meshes before rendering.

- **Open NURBS Surfaces**

 Control render faces with U/V DefResolu. Higher numbers give more faces. (Note normal errors).

- **Closed NURBS Surfaces**

 Control with DefResolu control. (Note the normal errors, and how implicit seam shows).

- **Curves and Text**

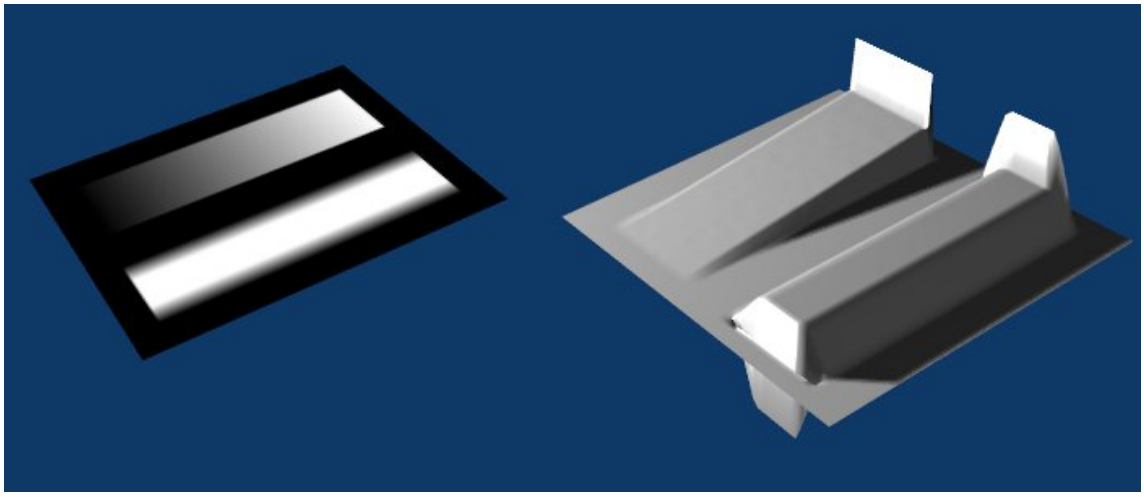
 Control with DefResolu control. Higher gives more render faces. (Note that the large flat surfaces have few renderfaces to displace).

Using the Displace Modifier

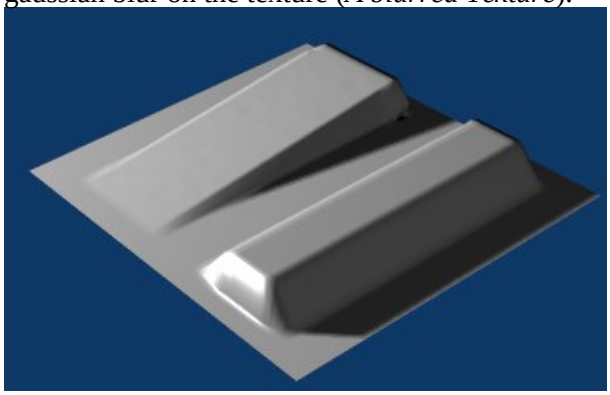
If you want more control over your displacement, you'll probably want to use the Displace Modifier. This feature has lots of different options so that you can customize the displacement excactly to your liking.

Examples

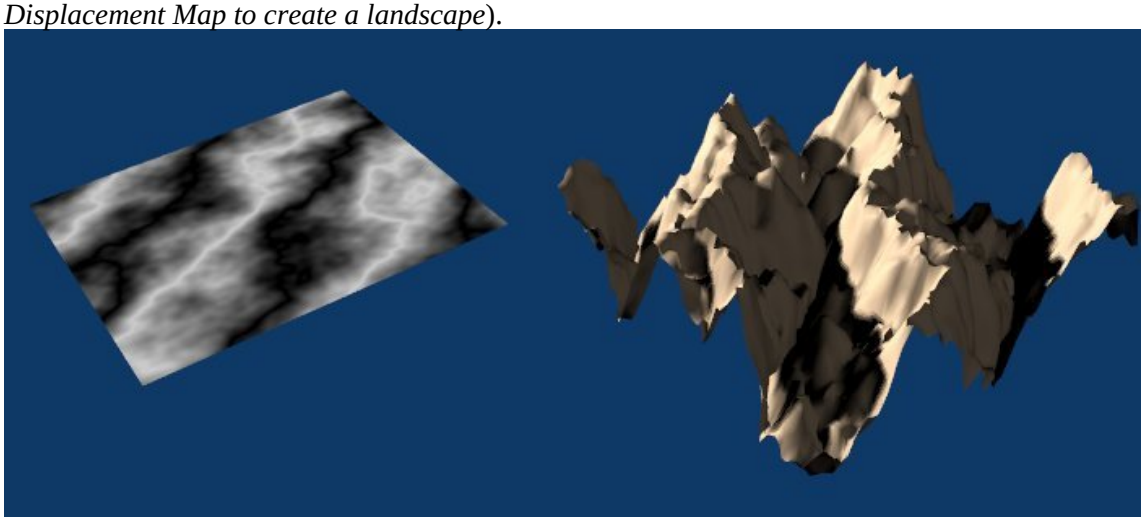
At first a not-so-well working example (*Texture and Displacement Map*):



Texture and Displacement Map. The result shows some errors.
The sharp contrasting transitions from black to white yield problems. To correct this, use a little bit of gaussian blur on the texture (*A blurred Texture*).



A blurred Texture yields the correct result.
If you use a Texture (like Marble) with no sharp transitions, the displacement works quite well (*A Displacement Map to create a landscape*).



A Displacement Map to create a landscape.
Advanced Materials often use Displacement Maps. Here a Marble Texture was applied to various Map To values, including Disp. The brink of the "comet" would be flat otherwise (*A Displacement Map for advanced materials*). The sphere has 1024 faces.



A Displacement Map for advanced materials.

See Also

- Bump and Normal Maps

How to create a Displacement Map

If you are using procedural textures, start with a flood of 50% gray as this neutral color will make a good base for your texture, as it does not cause any displacement. Some adjustment can be done using the Colors Panel Bright and Contr sliders if you desire. Continue shaping your map with blacks and whites with the Bright and Contr sliders. Done? Then back to the "Map To" panel (if you are not using a displace modifier) and check "Disp" and tweak the amount.

Sharp lines in Displacement Maps can cause normal problems, since a renderface can be requested to move only one of its verts a great distance relative to the other 2-3. You tend to get better results if a small gaussian blur is run on the image first.

Procedural Textures

Mode: All Modes

Panel: Shading/Textures Context

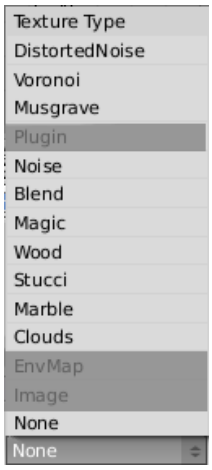
Hotkey: F6

Description

Procedural textures are textures that are defined mathematically. They are generally relatively simple to use, because they don't need to be mapped in a special way - which doesn't mean that procedural textures can't become very complex.

These types of textures are 'real' 3D. By that we mean that they fit together perfectly at the edges and continue to look like what they are meant to look like even when they are cut; as if a block of wood had really been cut in two. Procedural textures are not filtered or anti-aliased. This is hardly ever a problem: the user can easily keep the specified frequencies within acceptable limits.

Options



The Texture Type list in the Texture panel of the Texture Buttons (F6).
The non procedural textures are greyed out in *The Texture Type list*.

Nabla
: Almost all procedural textures in Blender use derivatives for calculating normals for texture mapping (with as exception Blend and Magic). This is important for Normal and Displacment Maps. The strength of the effect is controlled with the Nabla Number Button.

Hints

- Use the size buttons in the Map Input Panel to set the size that Procedural Textures are mapped to.
- Procedural textures can either produce colored textures, intensity only textures, textures with alpha values and normal textures. If intensity only ones are used the result is a black and white texture, which can be greately enhanced by the use of colorbands. If on the other hand you use colorbands and

need an intensity value, you have to switch on No RGB in the Map To panel.

Noise Basis

Mode: All Modes

Panel: Shading/Texture Context

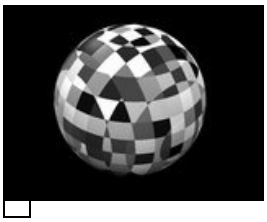
Hotkey: F6

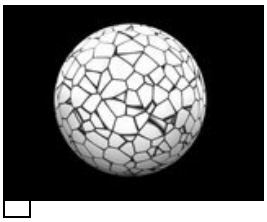
Description

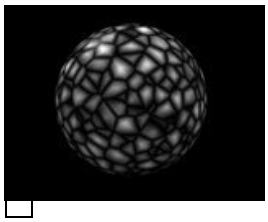
Each noise-based Blender texture (with the exception of Voronoi and simple noise) has a Noise Basis setting that allows the user to select which algorithm is used to generate the texture. This list includes the original Blender noise algorithm. The Noise Basis settings makes the procedural textures extremely flexible (especially Musgrave).

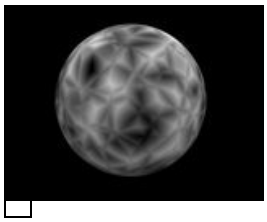
Examples

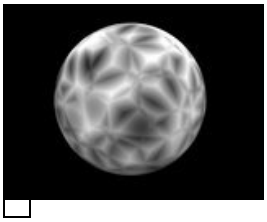
The Noise Basis governs the structural appearance of the texture.

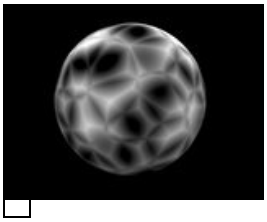

Cellnoise

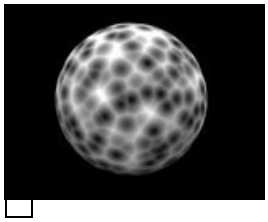

Voronoi Crackle

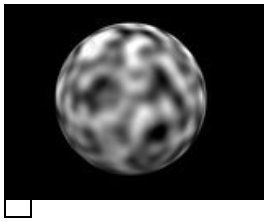

Voronoi F2-F1


Voronoi F4


Voronoi F3


Voronoi F2


Voronoi F1


Blender Original

There are two more possible settings for Noise Basis, which are relatively similar to Blender Original:

- Improved Perlin
- Original Perlin

Clouds

Mode: All Modes

Panel: Shading/Texture Context → Clouds

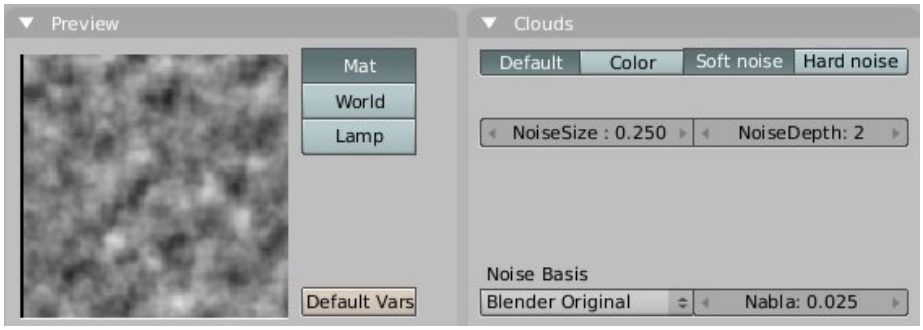
Hotkey: F6

Description

Often used for: Clouds, Fire, Smoke. Well suited to be used as Bumpmap, giving an overall irregularity to the material.

Result(s): Intensity (Default) or RGB-Color (Color)

Options



Clouds Texture Panels.

Default
: The standard Noise, gives an Intensity.

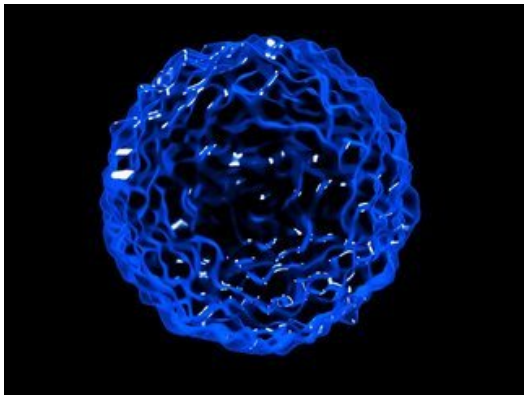
Color
: The Noise gives an RGB value.

Soft Noise/Hard Noise
: Changes the contrast and sharpness

NoiseSize
: The dimension of the Noise table.

NoiseDepth
: The depth of the Cloud calculation. A higher number results in a long calculation time, but also in finer details.

Examples



A Clouds Texture was used to displace the surface.

Technical Details

A three-dimensional table with pseudo random values is used, from which a fluent interpolation value can be calculated with each 3D coordinate (thanks to Ken Perlin for his masterful article "An Image Synthesizer", from the SIGGRAPH proceedings 1985). This calculation method is also called Perlin Noise. In addition, each noise-based Blender texture (with the exception of Voronoi and simple noise) has a new "Noise Basis" setting that allows the user to select which algorithm is used to generate the texture.

Marble

Mode: All Modes

Panel: Shading/Texture Context → Marble

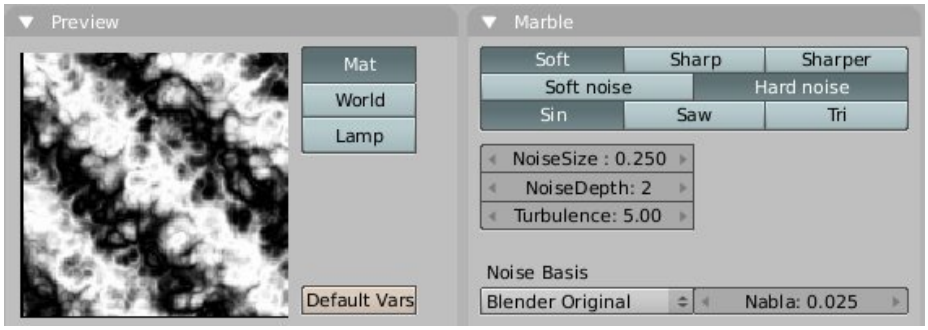
Hotkey: F6

Description

Often used for: Marble, Fire, Noise with a structure.

Result(s): Intensity value only.

Options



Marble Texture Panels

Soft/Sharp/Sharper
: Three pre-sets for soft to more clearly defined Marble.

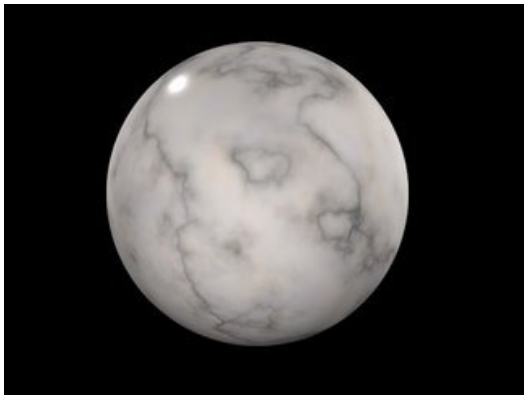
Soft Noise/Hard Noise
: The Noise function works with two methods.

NoiseSize
: The dimensions of the Noise table.

NoiseDepth
: The depth of the Marble calculation. A higher value results in greater calculation time, but also in finer details.

Turbulence
: The turbulence of the sine bands.

Examples



Marble from the Blender Materials Library

Technical Details

Bands are generated based on a sine formula and Noise turbulence.

Stucci

Mode: All Modes

Panel: Shading/Texture Context → Stucci

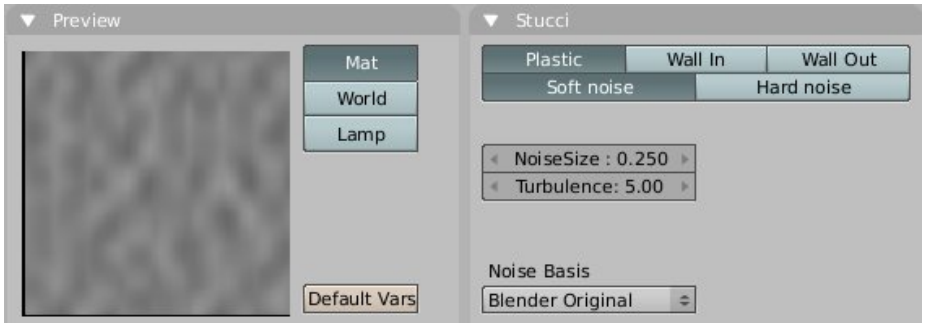
Hotkey: F6

Description

Often used for: Stone, Asphalt, Oranges. Normally for Bump-Mapping to create grainy surfaces.

Result(s): Normals and Intensity

Options



Stucci Texture Panels

Plastic
: The standard Stucci.

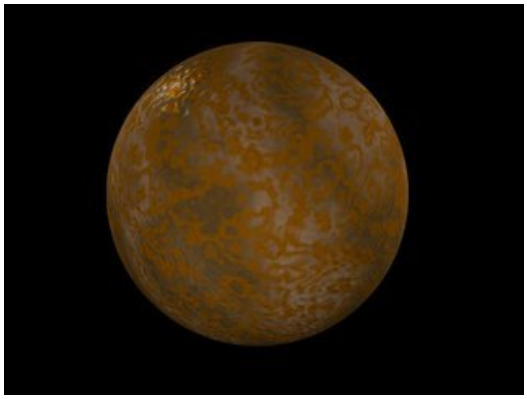
Wall In, Wall out
: This is where Stucci gets it name. This is a typical wall structure with holes or bumps.

Soft Noise/Hard Noise
: There are two methods available for working with Noise.

NoiseSize
: The dimension of the Noise table.

Turbulence
: The depth of the Stucci calculations.

Examples



Some rusty metal. Stucci was used to "Bump" the surface a bit.

Technical Details

Based on noise functions

Wood

Mode: All Modes

Panel: Shading/Texture Context → Wood

Hotkey: F6

Description

Often used for: Wood

Result(s): Intensity only

Options



Wood Texture Panels.

Bands
: The standard Wood texture.

Literal|Rings
: This suggests 'wood' rings.

BandNoise
: Gives the standard Wood texture a certain degree of turbulence.

RingNoise
: Gives the rings a certain degree of turbulence.

Soft Noise/Hard Noise
: There are two methods available for the Noise function.

NoiseSize
: The dimension of the Noise table.

Turbulence
: The turbulence of the BandNoise and RingNoise types.

Examples

See the section Doc:Tutorials/Textures/Wood for a method to create procedural wood.



"Wenge Wood" by Claas Eike Kuhnen.



Colorbands are used in both materials and textures, as well as other places where a range of colors can be computed and dipslayed. In the example to the right, we want to texture a snake, specifically the deadly coral snake. We want to make a repeating set of four colors: Black, yellow, red, yellow (and then back to black again). We also want to make the rings sharp in definition and transition. This example uses 8 color band settings: 0 and 7 are black; 1 and 2 are yellow, 3 and 4 are red, and 5 & 6 are yellow. Position 0 and 1 close together, 2 and 3, etc. Use a little noise and turbulence; together with the scales texture you should get really close!

Technical Details

Generation
: In this case, bands are generated based on a sine formula. You can also add a degree of turbulence with the Noise formula.

Coordinates
: As the band is based on a sine formula, the texture repeats itself every pi units rather than every 1.0 units. To correct this, scale the texture by a value of pi for the dimension you wish.

Magic

Mode: All Modes

Panel: Shading/Texture Context → Magic

Hotkey: F6

Description

Often used for: This is difficult, it was hard to find an application whatsoever. One could use it for "Thin Film Interference", if you set Map Input to Refl and use a relatively high Turbulence.

Result(s): RGB

Options

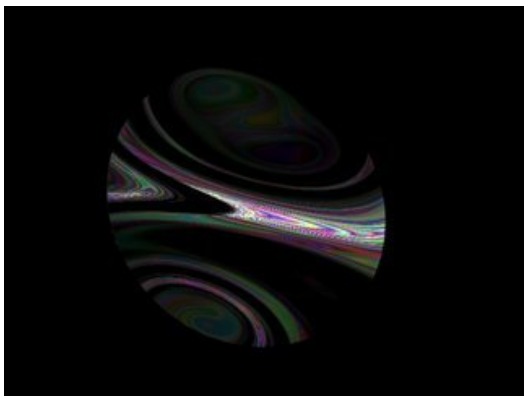


Magic Texture Panels.

Depth
: The depth of the calculation. A higher number results in a long calculation time, but also in finer details.

Turbulence
: The strength of the pattern.

Examples



"Thin Film Interference" with Magic Texture
I've used two Magic Textures in *"Thin Film Interference" with Magic Texture*. Both use the same texture with Depth 4, Turbulance 12. Both have Map Input set to Refl. The first texture is mapped to Nor, the second to Col.

Technical Details

The RGB components are generated independently with a sine formula.

Blend

Mode: All Modes

Panel: Shading/Texture Context → Blend

Hotkey: F6

Description

Often used for: This is one of the most important procedural textures. You can use blend textures to blend other textures together (with Stencil), or to create nice effects (especially with the Map Input: Nor trick). Just remember: if you use a colorband to create a custom blending, you may have to use No RGB, if the Map To value needs an intensity input!

Result(s): Intensity

Options



Blend Texture Panels

Lin
: A linear progression.

Quad
: A quadratic progression.

Ease
: A flowing, non-linear progression.

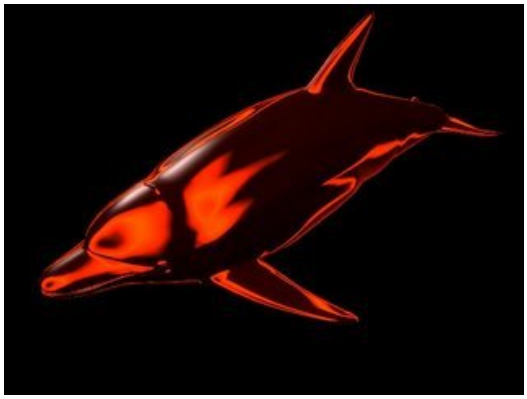
Diag
: A diagonal progression.

Sphere
: A progression with the shape of a three-dimensional ball.

Halo
: A quadratic progression with the shape of a three-dimensional ball.

Flip XY
: The direction of the progression is flipped a quarter turn.

Examples



A custom radial blend with Map Input set to Nor, Map To Ref and Emit. Blend file

Technical Details

The Blend texture generates a smoothly interpolated progression.

Noise

Mode: All Modes

Panel: Shading/Texture Context → Noise

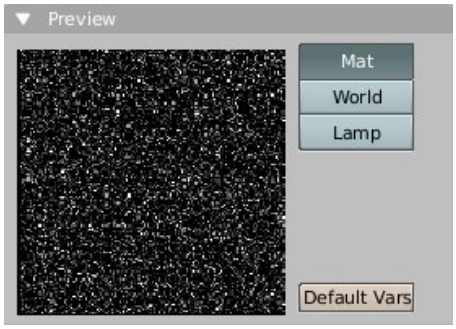
Hotkey: F6

Description

Often used for: White noise in an animation. This is not well suited if you don't want an animation. For material roughness take clouds instead.

Result(s): Intensity

Options



Noise Texture Panel.
There is no panel and no buttons. Just switch it on.

Technical Details

Although this looks great, it is not Perlin Noise! This is a true, randomly generated Noise. This gives a different result every time, for every frame, for every pixel.

Musgrave

Mode: All Modes

Panel: Shading/Texture Context → Musgrave

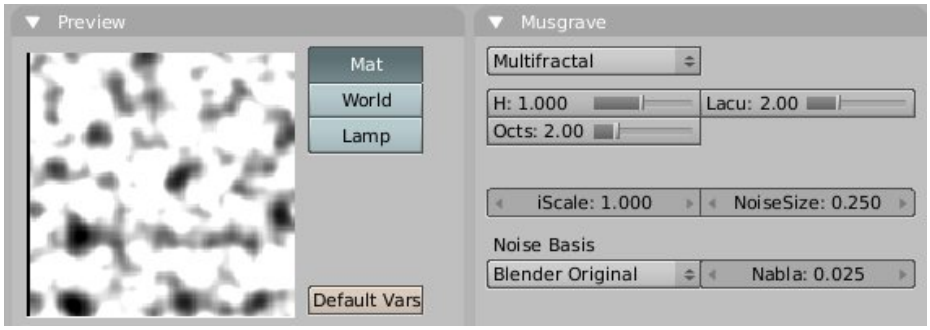
Hotkey: F6

Description

Often used for: Organic materials, but it's very flexible. You can do nearly everything with it.

Result(s): Intensity

Options



Musgrave Texture Panels.

Noise Types
: This procedural texture has five noise types on which the resulting pattern can be based and they are selectable from a dropdown menu at the top of the tab. The five types are:
 - ◊ fBm:
 - ◊ Hetero Terrain:
 - ◊ Hybrid Multifractal:
 - ◊ Ridged Multifractal:
 - ◊ Multifractal:

 These noise types determine the manner in which Blender layers successive copies of the same pattern on top of each other at varying contrasts and scales.

In addition to the five noise types, Musgrave has a noise basis setting which determines the algorithm that generates the noise itself. These are the same noise basis options found in the other procedural textures.

The main noise types have four characteristics which can be set in the number buttons below the dropdown list. They are:

H (Fractal Dimension) - Range 0 to 2)
: Fractal dimension controls the contrast of a layer relative to the previous layer in the texture. The higher the fractal dimension, the higher the contrast between each layer, and thus the more detail

shows in the texture.

Lacu (Lacuniarity) - Range 0 to 6)

Lacuniarity controls the scaling of each layer of the Musgrave texture, meaning that each additional layer will have a scale that is the inverse of the value which shows on the button. i.e. Lacu = 2 -> Scale = 1/2 original

Octs (Octave) - Range 0 to 8)

Octave controls the number of times the original noise pattern is overlayed on itself and scaled/contrasted with the fractal dimension and lacuniarity settings.

The Hybrid Multifractal, Ridged Multifractal, and Hetero Terrain types have additional settings:

Ofst (Fractal Offset)

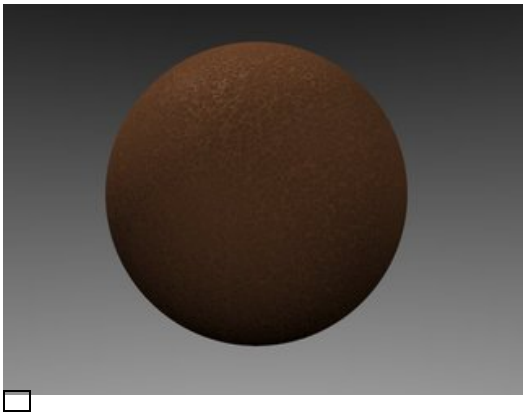
All three have a "Fractal Offset" button labeled Ofst. This serves as a "sea level" adjustment and indicates the base height of the resulting bump map. Bump values below this threshold will be returned as zero.

Gain

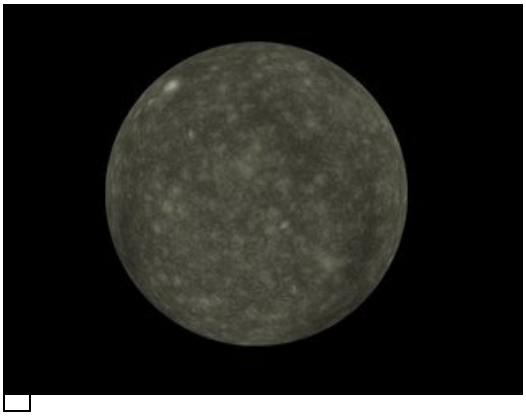
Hybrid Multifractal and Ridged Multifractal both have a Gain setting which determines the range of values created by the function. The higher the number, the greater the range. This is a fast way to bring out additional details in a texture where extremes are normally clipped off.

Examples

See the Samples Gallery from the release notes for more examples.



Leather made with a single Musgrave Texture



Stone made with a combination of 3 different Musgrave Textures.

Technical Details

More information about these textures can be found at the following URL: Musgrave Documentation

Voronoi

Mode: All Modes

Panel: Shading/Texture Context → Voronoi

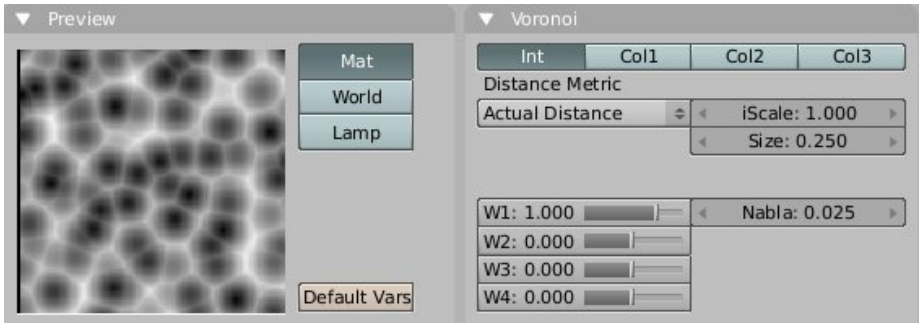
Hotkey: F6

Description

Often used for: Very convincing Metal, especially the "Hammered" effect. Organic shaders (e.g. scales, veins in skin).

Result(s): Intensity (default), Color.

Options



Voronoi Texture Panels.

Distance Metric
: This procedural texture has seven Distance Metric options. These determine the algorithm to find the distance between cells of the texture. These options are:
 - ◊ Minkovsky
 - ◊ Minkovsky 4
 - ◊ Minkovsky 1/2
 - ◊ Chebychev
 - ◊ Manhattan
 - ◊ Distance Squared
 - ◊ Actual Distance

 The Minkovsky setting has a user definable value (the Exp button) which determines the exponent (e) of the distance function $(x^e + y^e + z^e)^{1/e}$. A value of one produces the Manhattan distance metric, a value less than one produces stars (at **0.5**, it gives a Minkovsky 1/2), and higher values produce square cells (at **4.0**, it gives a Minkovsky 4, at **10.0**, a Chebychev). So nearly all Distance Settings are basically the same - variations of Minkowsky.

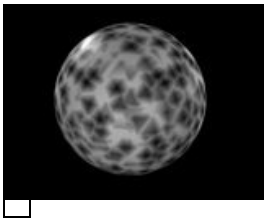
You can get irregularly-shaped rounded cells with the Actual Distance/Distance Squared options.

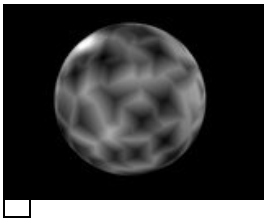
Four sliders at the bottom of the Voronoi panel represent the values of the four Worley constants (explained a bit in the Worley Documentation), which are used to calculate the distances between each cell in the texture based on the distance metric. Adjusting these values can have some interesting effects on the end result. Check the Samples Gallery for some examples of these settings and what textures they produce.

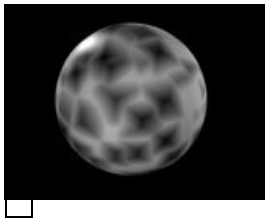
At the top of the panel there are four variation buttons which use four different noise basis as methods to calculate color and intensity of the texture output. This gives the Voronoi texture you create with the "Worley Sliders" a completely different appearance and is the equivalent of the noise basis setting found on the other textures.

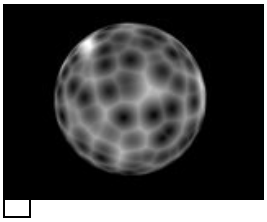
Examples

See the Samples Gallery from the release notes for more examples.

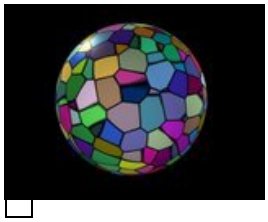

Manhattan

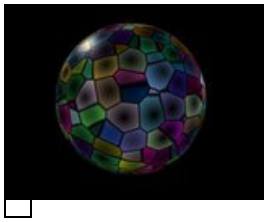

Cebychev


Minkowsky 10


Noise Basis Int.


Noise Basis Col1.


Noise Basis Col2.


Noise Basis Col3.

Technical Details

For a more in depth description of the Worley algorithm, see: Worley Documentation.

Distorted Noise

Mode: All Modes

Panel: Shading/Texture Context → Distorted Noise

Hotkey: F6

Description

Often used for: Grunge, very complex and versatile

Result(s): Intensity

Options



Distorted Noise Texture Panels.

Distortion Noise
: The texture to use to distort another

Noise Basis
: The texture to be distorted

Noise Size
: The size of the noise generated

Distortion Amount
: The amount that Distortion Noise affects Noise Basis

Examples

See the Samples Gallery from the release notes for more examples.

Technical Details

Distortion Noise takes the option that you pick from Noise Basis and filters it, to create hybrid pattern.

Image Textures

Mode: All Modes

Panel: Shading/Texture Context

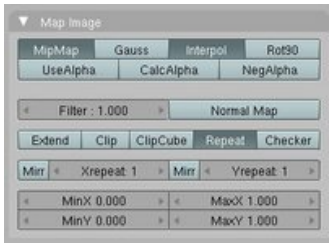
Hotkey: F6

Description

If you choose the Texture Type Image in the Texture Panel, the Map Image and Image panels appear, allowing you to control most aspects of image textures and how they are applied.

The term, "image texture," simply means that a graphic image (composed of R, G, B and perhaps Alpha information arranged in pixels) is used as the basic source of input to the texture. As with any other type of texture, this data can be used in any number of ways ... not just as, say, a "decal."

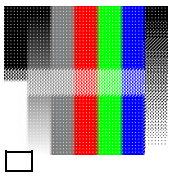
Map Image Panel Options



Map Image Panel in the Texture Buttons (F6).
This panel controls what aspects of the image are used, how it is to be mapped to the underlying material, and controls whether it is to be offset from its origin, and whether it should be repeated or stretched to fit.



Background image



Foreground image

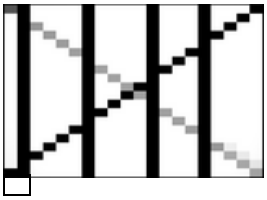
These two different images are used here to demonstrate the different Map Image options. *The background image* is an ordinary JPG-file, *the foreground image* is a PNG-file with various Alpha- and Greyscale values. The vertical bar on the right side of the foreground image is an Alpha blend, the horizontal bar has 50% Alpha.}}

MipMap
: MipMaps are precalculated, smaller, filtered Textures for a certain size. A series of pictures is generated, each half the size of the former one. This optimizes the filtering process. By default, this option is enabled and speeds up rendering (especially useful in the game engine). When this option is OFF, you generally get a sharper image, but this can significantly increase calculation time if the filter dimension (see below) becomes large. Without MipMaps you may get varying pictures from slightly different camera angles, when the Textures become very small. This would be noticeable in an animation.

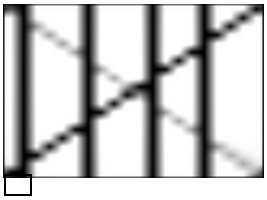
Gauss

Used in conjunction with MipMap, it enables the MipMap to be made smaller based on color similarities. In the game engine, you want your textures, especially your MipMap textures to be as small as possible to increase rendering speed and framerate.

InterPol



Enlarged Imagetexture without InterPol.



Enlarged Imagetexture with InterPol.

This option interpolates the pixels of an Image. This becomes visible when you enlarge the picture. By default, this option is on. Turn this option OFF to keep the individual pixels visible and if they are correctly anti-aliased. This last feature is useful for regular patterns, such as lines and tiles; they remain 'sharp' even when enlarged considerably. When you enlarge this 10x10 pixel Image , the difference with and without InterPol is clearly visible (*Enlarged Imagetexture*). Turn this image off if you are using digital photos to preserve crispness.

Rot90

Rotates the Image 90 degrees counterclockwise when rendered.

UseAlpha



Foreground Image with UseAlpha. The alpha values of the pixels are evaluated.



Foreground Image with CalcAlpha.

Works with PNG and TGA files since they can save transparency information (*Foreground Image with UseAlpha*). Where the alpha value in the image is less than 1.0, the object will be partially transparent and stuff behind it will show.

CalcAlpha

Calculate an alpha based on the RGB values of the Image. Black (0,0,0) is transparent, white (1,1,1) opaque (*Foreground Image with CalcAlpha*). Enable this option if the image texture is a mask. Note that mask images can use shades of gray that translate to semi-transparency, like ghosts, flames, and smoke/fog.

NegAlpha

Reverses the alpha value. Use this option if the mask image has white where you want it transparent and vice-versa.

Filter

The filter size used in rendering, and also by the options MipMap and Interpol. If you notice gray lines or outlines around the textured object, particularly where the image is transparent, turn this value down from 1.0 to 0.1 or so.

Normal Map

This tells Blender that the image is to be used to create the illusion of a bumpy surface, with each of the three RGB channels controlling how to fake a shadow from a surface irregularity. Needs specially prepared input pictures. See section Bump and Normal Maps.

Usually, an image texture will fit the entire size of the texture space. Similarly to the NoiseSize options in procedural textures, you can also change the scaling and positioning of image textures, within the texture

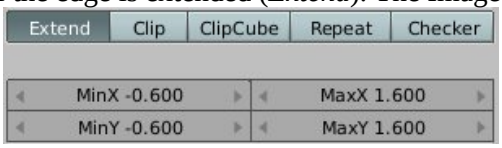
itself, before it is mapped.

Extend



Extend. Outside the Image the colour of the edge is extended.

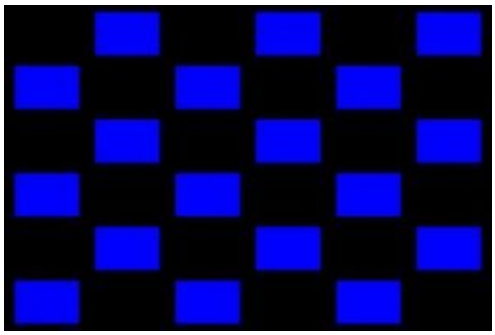
Outside the Image the colour of the edge is extended (*Extend*). The Image was put in the center of the object with the shown settings.



Clip
: Outside the Image, an alpha value of 0.0 is returned. This allows you to 'paste' a small logo on a large object.

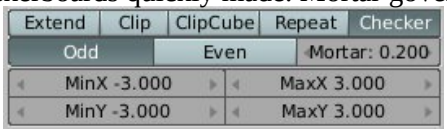
ClipCube
: The same as Clip, but now the 'Z' coordinate is calculated as well. Outside a cube-shaped area around the Image, an alpha value of 0.0 is returned.

Checker



Checker generates Checkerboard patterns. Here a blue picture was used on a black background.

Checkerboards quickly made. Mortar governs the distance between the checkers in parts of the texture size.



(You can use the option size on the Map Input Panel as well to create the desired number of checkers).

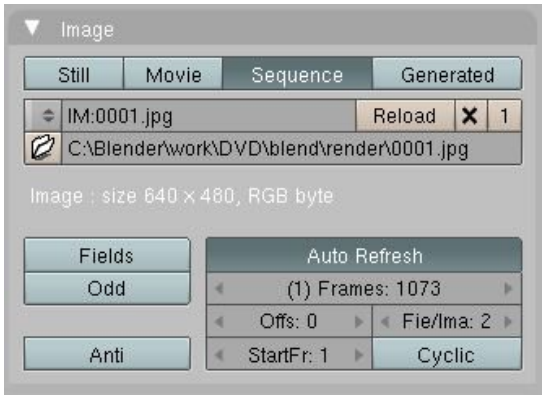
Mirror

The two Mirr buttons allow you to map the texture as a mirror, or automatic flip of the image, in the corresponding X and/or Y direction.

Repeat
: The Image is repeated horizontally and vertically as often as set in Xrepeat and Yrepeat.

MinX, MinY, MaxX, MaxY
: The offset and the size of the texture in relation to the texture space. Pixels outside this space are ignored. Use these to crop, or choose a portion of a larger image to use as the texture.

Image Panel Options



The Image Panel in the Texture Buttons (F6).

Load Image

- Load a single image file in one of Blender's supported file formats: BMP, JPG, PNG, TGA, TIFF, OpenEXR, Cineon, DPX and Radiance HDR. Others, like PSD and GIF - are partially supported via QuickTime on Windows and Mac versions.
- To use a vector format SVG image, use the [Vectrex texture plugin]
- Load an animation as a numbered image sequences in any of the supported image file formats. To do this, first click on the first frame and then Load file. Then change the Image type to Sequence, and enter the ending frame number
- Load an animation as an avi or mov file. Accepted are uncompressed or JPG compressed AVIs. On Windows and Mac with QuickTime, most quicktime movies should work, and on Windows platforms all videos with a supported codec should work.

 For any of the above, you can use absolute as well as Relative Paths. The // (double Backslash) means the working directory, .. (two dots) point at the parent directory. Select relative or absolute in the file browser window header.

Movie
: Movie files (AVIs supported by Blender, SGI-movies) and "anim5" files can also be used for an Image texture. In this case the subsequent Panel, Anim and Movie is populated by Buttons.

IM
: The internal name of the IMage. Use the select button to rapidly change between textures loaded into memory. â § ShiftLMB click into the field to give the image a meaningful name

Reload
: Load the image again from disk (i.e. if it's been changed in an external application)

X

Delete the link from this texture to this image

Users
The number of other materials that use this image. Click to make the information about this image local to this instance of use.

Pack ()
Saves the still or generated image inside the blend-file. Use this if you are mailing or sharing .blend files between PC's with different directory structures.

The text below the filename provides information about the file

Fields
Video frames consist of two different images (fields) that are merged by horizontal line. This option makes it possible to work with field images. It ensures that when Fields are rendered the correct field of the Image is used in the correct field of the rendering. MipMapping cannot be combined with Fields.

Odd
Normally, the first field in a fielded (interlaced) video frame begins on the first line. Some frame grabbers do this differently.

Anti
Graphic images such as cartoons and pictures that consist of only a few colors with a large surface filling can be anti-aliased as a built in pre-process. This is not useful for Photos and similar pictures. Needs OSA.

Auto Refresh
When you change frames, blender gets the latest image from the sequence or codec.

Frames
How many frames of the animation to use; the length of the segment

Offs
What frame number inside the movie/sequence to start grabbing

Fie/Ima
Fields per Image. Used with Fields and interlaced video, it says whether each image has both odd and even, or just one.

StaFr
Starting Frame - when the animation reaches this frame number, the movie/sequence will start playing

Cyclic
When the video ends, it will loop around the to the start and begin playing again.

Seamless Images



The good, the bad, and the ugly
When an image is repeated, the left side of the repeat copy is mushed up against the right side. Similarly, when repeating upwards, the bottom of the second copy is placed against the top of the first. Where they meet, a seam may show if the pixels of one side do not line up with the other. You want image textures to be

seamless, so that when viewed side by side, the repeating is not detectable. The best example of this is a brick texture. Use an image texture of a few bricks where the boundaries of the image align at the top and bottom on the mortar, and the left and right edges evenly split a brick, and that the left brick matches up nicely with the far right brick. This concept applies to terrain (outside grass, sidewalks, pavement), flooring (inside wood, stone, tile), tree bark, animal skins, roofing shingles of all types, cloth and fabric designs, wallpaper designs, dirty surfaces, sidewalks, and almost all exterior building surfaces (continuous, siding, shingle, brick, stone veneer). Both images to the right were tiled 2 times across and 3 times vertically.

When working with existing images to make them seamless, try to correct any skewing based on camera angle; you want an image taken "straight on", because the program will repeat the image flat onto the surface. Also pay attention to lighting; you want a flat light across the surface, because if the left side is lighter than the right, when they are placed together, there will be a noticeable difference.

Some paint programs have features to help you. For example, GIMP has a map filter to Make Seamless any image. It also has tools like Threshold that show you any lighting differences.

Also, placing a prop in front of a seamed wall really fakes out the eye from detecting a pattern. A simple table and lamp or chair for inside walls, or a plant or person or car for outside really helps deflect attention from percieving a repeating pattern.

Video Textures

Mode: All Modes

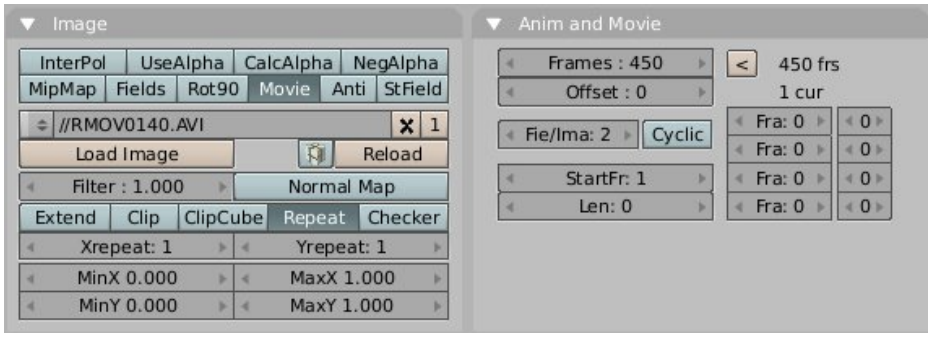
Panel: Shading/Texture Context → Anim and Movie

Hotkey: F6

Description

As well as animating a texture's mapping with its associated Ipo curves (eg. the texture's offset - ofsX/Y/Z), you can also use animated image sources as textures. The simplest method to get an animated texture is to use a video file. The video needs the same - or an integral division - number of frames per second (FPS) as the animation, to run at the same speed.

Options



An AVI file as an Image Texture.

frs / <
: The number frs in the Anim and Movie Panel shows how many frames were recognised. You can copy this number to the field Frames with the arrowbutton. The number cur shows, which frame is shown in the Preview Panel.

Frames
: This activates the animation option; another image file (in the same Image block) will be read per rendered frame (see also Fie/Ima). The number in the field Frames is the number of frames that shall be used in the animation. The last frame will be used as static texture for the rest of the animation, if you don't turn on the option Cyclic.

Offset
: The number of the first picture of the animation. The end frame is calculated as Frames + Offset.

Fie/Ima
: "Fields per Image:" The number of fields per rendered frame. If no fields are rendered, even numbers must be entered here. (2 fields = 1 frame). This sets the speed of the animation. The correct settings depend on the framerate of the texturevideo, the framerate of the rendered animation, whether you render Fields (in the Render Panel of the Szene context) and whether the texturevideo uses "Fields" (Fields Button in the Image Panel of the Texture buttons). Some examples:
 - ◊ The video has 24 FpS, the animation shall have 24 FpS. You are rendering without Fields. Set Fie/Ima to 2.
 - ◊ The video has 12 FpS, the animation shall have 24 FpS. You are rendering without Fields. Set Fie/Ima to 4.

◊ The video has 16 Frames, the animation shall have 96 Frames. You are rendering without Fields. Set Fie/Ima to 6.
◊ The video has 24 FpS, the animation shall have 24 FpS. You are rendering with Fields and are using Fields in the Image Panel. Set Fie/Ima to 1.

Cyclic
The animation Image is repeated cyclically.

StartFr
The moment - in Blender frames - at which the animation Image must start. Until that moment the first Image of the video is used as texture.

Len
This button determines the length of the animation. A Len of 0 means, that the length is equal to Frames. By assigning Len a higher value than Frames, you can create a still at the end of the animation when you use cyclic.

Fra
The Fra buttons allow you to create a simple montage within an animation Image. The left button, Fra indicates the frame number, the right-hand button indicates how long the frame must be displayed (Stutter-Mode). An example follows. If you use Fra you have to set Frames and Len accordingly.

Brightness
To keep the original brightness of the video, use the option **Shadeless** for the material.

Numbered Image Sequences

Instead of a video file you can also use a numbered image sequence. The simplest procedure will be to save the images in a subdirectory to your blend file and to load **one** of the images from this directory.

Examples

This example uses an image sequence instead of a video - 12 Image files (01.jpg to 12.jpg). The entry in the field Frames activates the animation. Now Blender tries to find the next frames by changing a number in the filename. You may not use the option Movie also!

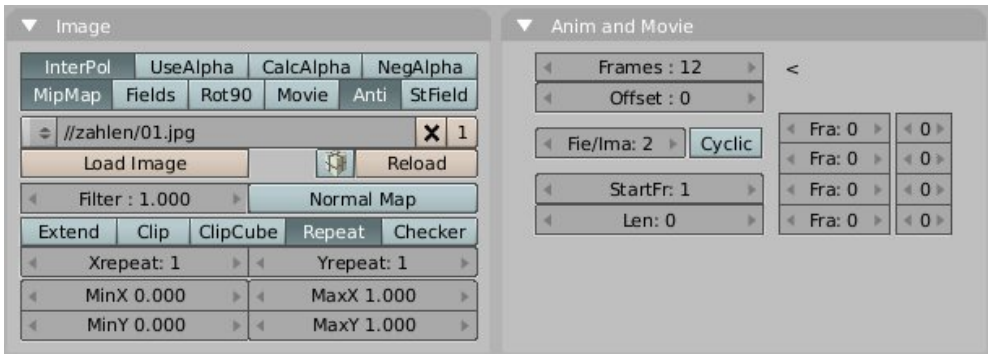
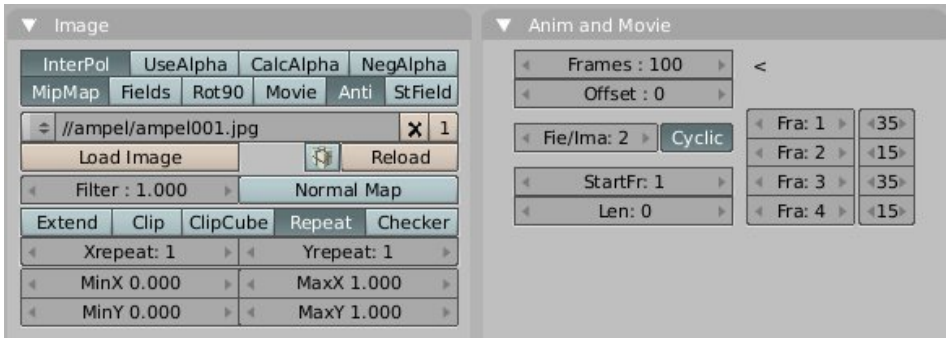
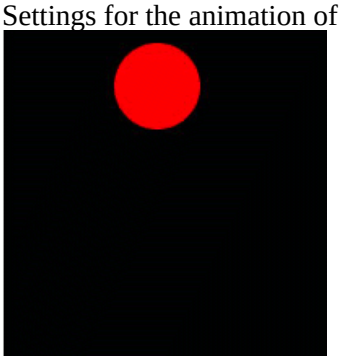


Abbildung 2: Nummerierte Bilddateien als Textur
Everything else works similarly to video textures. Here's an example of the Fra Option.
Lets assume, you would like to create an animated Traffic light. At first you create four different images: Red, Red/Yellow, Green, Yellow. These four images shall change continuously.



Settings for the animation of a Traffic light with four images. Idea by "tordat" [1]



The Traffic light with the settings from the picture *Settings for the animation*.
Please note that you need 100 Frames, though only four images are used. The result is not surprising (*The Traffic light*).

Limitations

When using image sequences, your image sequence must be named a certain way, to keep Blender counting:

1. Blender tries to find the other files by changing a number in the file name. Only the rightmost digit is interpreted for this. For example: 01.ima.099.tga + 1 becomes 01.ima.100.tga.
2. The numbers have to have the samge length. Amend leading zeros. Blender counts from 1 to 12, but not back to 1. Instead it begins at 10. So use 01, 02 ...

Environment Maps

Mode: All Modes

Panel: Shading/Texture Context → Envmap

Hotkey: F6

Description

Environment maps take a render of the 3D scene and apply it to a texture, to use for faking reflections. If you want to achieve a very realistic result, raytraced reflections are a good solution. Environment Maps are another way to create reflective surfaces, but they are not so simple to set up.

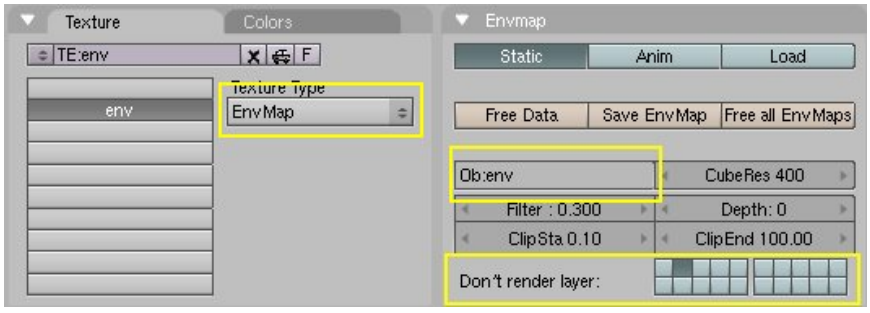
So why should one use Environment Maps?

- The main reason is probably that they can be much faster than raytracing reflections. In certain situations they need to be calculated only once, and may be reused like any ordinary texture. You may even modify the precalculated Environment Map in an image editor.
- Environment maps can also be blurred and render even faster because the resolution can then be lowered. Blurring a reflection with the raytracer always adds to the render time, sometimes quite a lot.
- Halos (a visualization type for particles) are not visible to raytraced reflections, so you need to setup environment maps to reflect them.
- Keypoint strands (another visualization type for particles) are also not visible to raytraced reflections, so you need to setup environment maps to reflect them.

Just as we render the light that reaches the viewing plane using the camera to define a viewpoint, we can render the light that reaches the surface of an object (and hence, the light that might ultimately be reflected to the camera). Blender's environment mapping renders a cubic image map of the scene in the six cardinal directions from any point. When the six tiles of the image are mapped onto an object using the Refl input coordinates, they create the visual complexity that the eye expects to see from shiny reflections.

Note
It's useful to remember here that the true goal of this technique is *believability*, not *accuracy*. The eye doesn't need a physically accurate simulation of the light's travel; it just needs to be lulled into believing that the scene is real by seeing the complexity it expects. The most unbelievable thing about most rendered images is the sterility, not the inaccuracy.

Options



Reflecting plane EnvMap settings.

Important
For correct results, the mapping of an environment map texture must be set to 'Refl' (reflection co-ordinates) in the Map Input panel of the Material context.

Ob
: Environment maps are created from the perspective of a specified object. The location of this object will determine how 'correct' the reflection looks, though different locations are needed for different reflecting surfaces. Usually, an Empty is used as this object.

- ◊ For planar reflections, the object should be in a location mirrored from the camera, on the other side of the plane of reflection (see Examples). This is the most accurate usage of Environment maps.
- ◊ For spherical reflections, the object should be in the center of the sphere. Generally, if the reflecting sphere's object center point is in the center of its vertices, you can just use the name of the actual sphere object as the Ob:
- ◊ For irregular reflections, there's no hard and fast rule, you will probably need to experiment and hope that the inaccuracy doesn't matter.

Don't render layer
: The layer to exclude from the environment map creation. Since environment maps work by rendering the scene from the location of the Ob: object, you will need to exclude the actual reflecting surface from the environment map, otherwise it will occlude other objects that should be reflected on the surface itself.

 Eg. If you are rendering an environment map from the center of a sphere, all the environment map will show by default is the inside of the sphere. You will need to move the sphere to a separate layer, then exclude that layer from the environment map render, so that the environment map will show (and hence reflect) all the objects outside the sphere.

CubeRes
: The resolution of the cubic environment map render. Higher resolutions will give a sharper texture (reflection), but will be slower to render.

Filter
: The amount of blurring applied to the texture. Higher values will blur the environment map to fake blurry reflections.

Depth
: The number of recursive environment map renders. If there are multiple reflecting objects using environment maps in the scene, some may appear solid, as they won't render each other's reflections. In order to show reflections within reflections, the environment maps need to be made multiple times, recursively, so that the effects of one environment map can be seen in another environment map. See Examples.

Clipsta/ClipEnd
: The clipping boundaries of the virtual camera when rendering the environment map

Blender allows three types of environment maps, as you can see in *Reflecting plane EnvMap settings.*:

Static
: The map is only calculated once during an animation or after loading a file.

Anim

The map is calculated each time a rendering takes place. This means moving Objects are displayed correctly in mirroring surfaces.

Load
When saved as an image file, environment maps can be loaded from disk. This option allows the fastest rendering with environment maps, and also gives the ability to modify or use the environment map in an external application.

When using planar reflections, if the camera is the only moving object and you have a reflecting plane, the Empty must move too and you must use Anim environment map. If the reflecting object is small and the Empty is in its center, the environment map can be Static, even if the object itself rotates since the Empty does not move. If, on the other hand, the Object translates the Empty should follow it and the environment map be of Anim type.

Free Data
Clears the currently rendered environment map from memory. This is useful to refresh a Static environment maps and you have changed things in your scene since the last time the environment map was rendered. Anim environment maps do this automatically on every render.

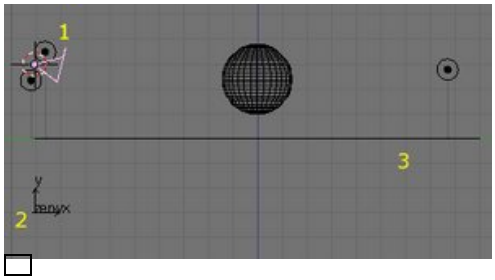
Save EnvMap
Saves the currently stored static environment map to disk as an image file. This can be loaded again with Load.

Free all EnvMaps
Does the same as Free Data, but with all environment maps in the scene. This is a useful shortcut when using recursive environment maps (when the Depth is greater than 0).

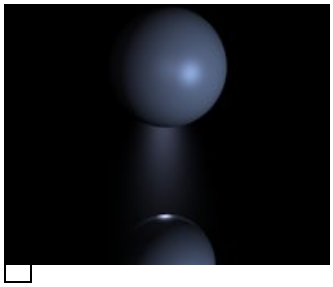
Note
EnvMap calculation can be disabled at a global level by the EnvMap Tog Button in the Render Panel of the Rendering Buttons.

Examples

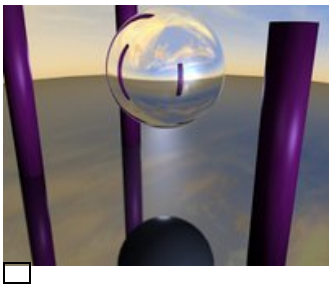
In this example, an empty is used as the Ob: of the reflecting plane's environment map. It is located in the specular position of the camera with respect to the reflecting surface. (This is possible, strictly speaking, only for planar reflecting surfaces.) Ideally, the location of the empty would mirror the location of the camera across the plane of the polygon onto which it is being mapped.



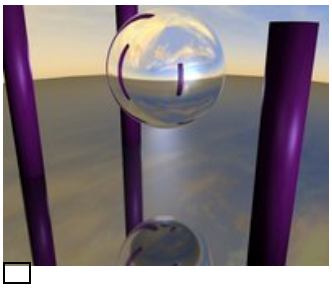
Planar reflection example. 1: Camera, 2: Empty, 3: Reflecting Plane.


Sphere on a reflecting surface.

The following images show the effect of the Depth. The first render has depth set to 0. This means the environment map on the plane has rendered before the environment map of the sphere, so the sphere's reflection isn't shown. By raising the Depth, the environment map is rendered recursively, in order to get reflections of reflections.



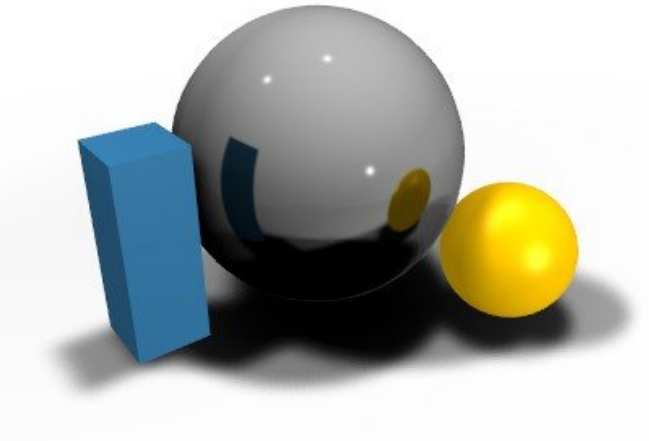
Reflecting sphere on a reflecting surface.



Reflecting sphere on a reflecting surface with multiple reflections.

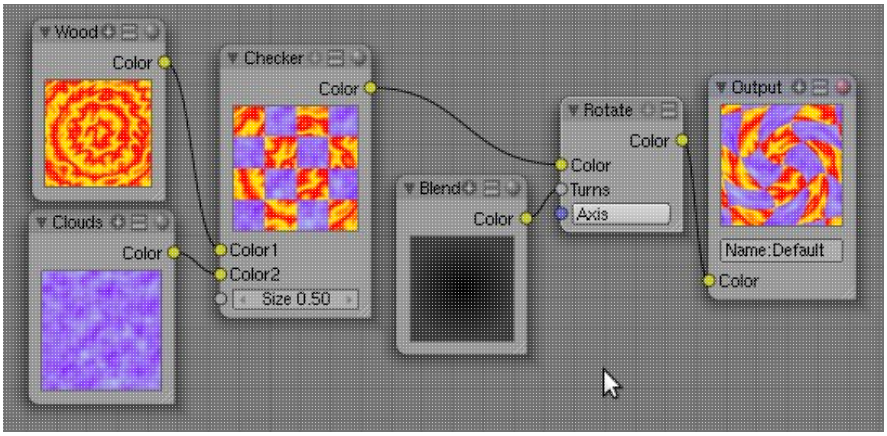
Limitations

Because environment maps are calculated from the exact location of the Ob:'s object center, and not from actual reflecting surface, they can often be inaccurate, especially with spheres. In the following image, the rectangular prism and the smaller spheres are touching the sides of the large reflecting sphere, but because the environment map is calculated from the center of the sphere, the surrounding objects look artificially far away.



Inaccurate spherical reflection, the coloured objects are artificially offset

Introduction



A texture node system example.

As an alternative to using the Texture Channels, Blender includes a node-based texture generation system which enables you to create textures by combining colors, patterns and other textures in much the same way

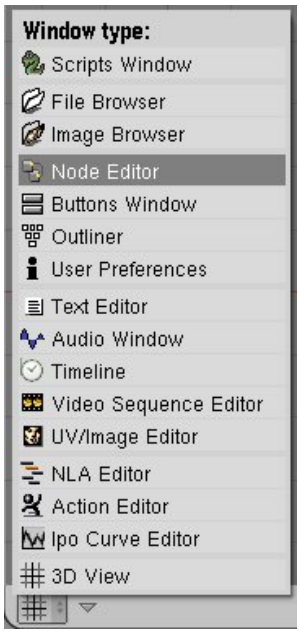
that you combine Material Nodes.

You can use these textures wherever you can use regular textures: you can place them in texture channels, in material nodes, in particle systems, and even inside other textures.

As the above illustration suggests, node networks ("noodles") contain three general types of nodes: **input** nodes, **filter** (or **transformation**) nodes, and **output** nodes. You can include any number of these nodes into the network, and connect them in any number of ways. This gives you limitless creative and technical control.

Note
Node-based textures do **not** work for realtime display, they will only be visible in rendered images.

Using Texture Nodes

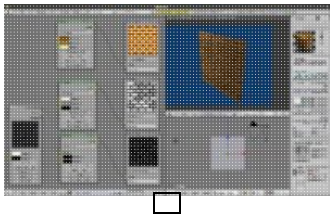


Opening a Node Editor.
To use texture nodes with the current texture, open a Node Editor window, set it to Texture mode by clicking the â← Textureâ← icon in its header (), and then click Use Nodes.

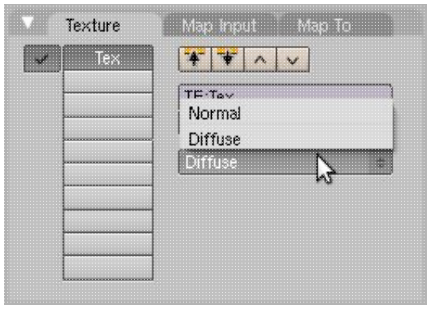
The default node setup will appear: a red-and-white checkerboard node connected to an Output named â← `Default`â← . For *texture* nodes, you can create as many Outputs as you like in your node setup. (Other types of node networks, as you may recall, are limited to only one Output node.) See the next section for details.

For instructions on how to add, remove and manipulate the nodes in the tree, see the Node Editor manual.

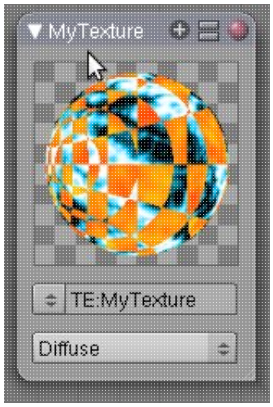
Using Multiple Outputs



An example multi-output tree.



Selecting an output using Texture Channels.



Selecting an output using Material Nodes.

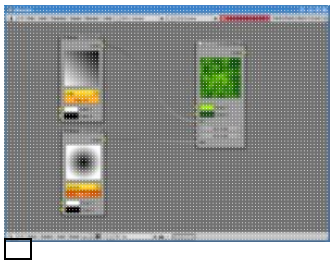
Each texture that you define with Texture Nodes can have several outputs, which you can then use for different things. For example, you might want your texture to define both a diffuse (color) map and a normal map. To do this, you would:

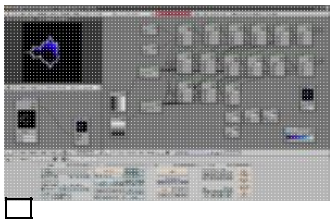
1. Add two Output nodes to the tree, and type new names into their Name text-boxes: *e.g.* â← `Diffuse`â← for one and â← `Normal`â← for the other.
2. In the Material sub-context for the material you want to use with the texture, add two Texture Channels. Map one to diffuse (Col), and the other to Normal. Set them both to your new texture.
3. Underneath the texture picker in the texture channels panel, youâ← ll see a dropdown list with the names of your outputs. In the channel you mapped to Diffuse, pick â← `Diffuse`â← , and in the channel you mapped to Normal, pick â← `Normal`â← .

You can also use these named outputs if you've decided to define your material using Material Nodes. In this case, you probably won't be using Texture Channels. Instead, you'll insert Texture nodes into your Material Node tree using Add â Input â Texture. Then, inside the texture node that you've just added, you can select which output you want to use *(e.g.* `Diffuse` or `Normal`).

Setup examples







See also

Development page

Texture Plugins

Mode: All Modes

Panel: Shading/Texture Context → Plugin

Hotkey: F6

Description

Texture plugins are external files that can be loaded up in the Blender interface, which provide controls in the Plugin panel based on what they can do. A plugin texture is a dynamically loaded library that exists as a separate file on your computer. When called in, it generates the texture. A standard set of plugin textures are distributed with Blender and are located in the `Blender Foundation/Blender/plugins/texture` directory, or wherever the pathspec points to in the User Preferences â File Paths â Tex Plugins field.

These plugins are developed by various people, and a general central collection site is hosted at the Blender Plugin Repository. See also this page for plugins lists and links. A recent brick texture can be found through the BlenderArtists forum. When you find a good texture plugin, we recommend that you put a copy in your personal library (lib/plugin/texture) directory.

Blender 2.46+
Plugins now work correctly with multi-threaded rendering. Prior versions using multi-threads will result in black or white spots or streaks in the render. To work around this, use single-threaded rendering.

Sample List of Available Texture Plugins

For the current list, examples, and more information about each plugin, please click the above link. Just to give you a quick idea as to the diversity and range of textures that are available as of January 2008:

afterglow
: Afterglow unlike Other plugins is both a Texture and a Sequence plugin for glowing objects in Blender.

brick
: Brick-drawing plugin, based on a sample plugin from NaN.

brick1
: Creates a brick texture.

ceramictiles
: This is a C implementation of the BMRT ceramictiles shader written by Larry Gritz.

chip
: This code is based at Ken Perlin's famous works...

chip2
: This code is based at Ken Perlin's famous works...

circdots_rgb
: blender plugin for generating circular dots (source written/compiled by sylvio sell, 2006)

clouds2
: Creates a nice cloud texture similar to the builtin one.

dots2
: Makes nice pok-a-dots.

fixnoise
: Creates a random static noise map, use this if your doing an animation instead of the builtin noise.

greenmarble

This is a C implementation of the BMRT greenmarble shader written by Larry Gritz.

led
: Creates LED numbers like on a digital clock.

lyapunov
: blender plugin for generating Lyapunov fat fractals. (source written/compiled by sylvio sell, dec 2004)

mandeltex
: blender plugin for exploring Benoit Mandelbrot's fractal set. (source written/compiled by rob haarsma, between dec 95 and april 1999)

matrix
: Creates noise similar to clouds and novichip.

musgrave
: a procedural texture plugin for blender that generates (and bumps !) Kent Musgrave's fractal noise patterns.

novichip
: Another noise plugin similar to clouds or matrix

pattern
: This plugin uses the sin function to create various patterns Similar to the sinus plugin a bit different however.

pie
: blender plugin for pie shaped divs.

r_weave
: Weave plugin that does fake bumpmapping and is tileable.

refract
: Fakes Refraction.

rings2
: This produces a ring-like pattern, much like the wood texture, but this is a 4D texture. (OK, 3D really, but the third is time, not Z). acceleration.

rtilings
: with this plugin you can create regular tilings of a plane (based on squares, triangles or hexagons).

sarah0
: The sarah0 plugin (a.k.a. Sarah's first plugin) If you like this plugin you should really check out the main website for the latest version. Bricks, tiles, interlocking patterns, you name it.

scales
: This plugin generated a texture pattern that resembles fish scales. Through settings, I suppose it could also do a round shingled roof.

sinus
: This plugin uses the sin function and some constants you can modify to get different textures

spirals
: Spiral Drawing plugin.

t_bricks
: This plugin generates brick textures with bumpmapping.

t_clouds
: This plugin generates cloud textures with bumpmapping.

t_marble
: This plugin generates marble and or wood textures with bumpmapping.

t_marble_terrain
: This plugin generates marble textures with slope and altitude constraints and bumpmapping.

t_terrain
: This plugin generates cloud textures with slope and altitude constraints and bumpmapping.

t_wood
: This plugin generates wood textures with bumpmapping.

tiles
: Creates a nice checkerboard pattern you can also add some noise to it.

trellis
: Create various regular 2D tiling patterns, also a true 3D cubic texture (like cells in POVRAY)

voronoi
: a procedural texture plugin for blender that generates Voronoi cell structures.

water
: Animate various kinds of ripples of water, like rain or just a dripping tap. There are many settings and they can be distorted by 'wobbly' noise. Very nice for all kinds of "wet" effects.

wbricks
: more advanced brick texture

Options

Load Plugin
: Opens a file select window to browse for a plugin to load. These plugins are **dll** files on Windows, **.so** files on Mac and various Unix flavors.

Once a plugin is loaded it turns the Texture Buttons window into its own set of buttons, as described in the individual plugin references.

Technical Details

Blender allows the dynamic linking at run time of shared objects, both texture and sequence plugins. In both cases these objects are pieces of C code written according to a given standard (*chapter_plugin_reference*). In the case of texture plugins, these chunks of code define functions accepting coordinates as input and providing a Color, Normal and Intensity output, exactly as the procedural Textures do.

See Also

- Blender's Plugin System

The most flexible way of mapping a 2D texture over a 3D object is a process called "UV mapping". In this process, you take your three-dimensional (X,Y & Z) mesh and unwrap it to a flat two-dimensional (X & Y ... or rather, as we shall soon see, "U & V") image. Colors in the image are thus mapped to your mesh, and show up as the color of the faces of the mesh. Use UV texturing to provide realism to your objects that procedural materials and textures cannot do, and better details than Vertex Painting can provide.

UV Explained



Box being inspected



Box mapped flat

The best analogy to understanding UV mapping is cutting up a cardboard box. The box is a three-dimensional (3D) object, just like the mesh cube you add to your scene.

If you were to take a pair of scissors and cut a seam or fold of the box, you would be able to lay it flat on a tabletop. As you are looking down at the box on the table, we could say that U is the left-right direction, is V is the up-down direction. This image is thus in two dimensions (2D). We use **U** and **V** to refer to these "texture-space coordinates" instead of the normal **X** and **Y**, which are always used (along with **Z**) to refer to "3D space."

When the box is reassembled, a certain UV location on the paper is transferred to an (X,Y,Z) location on the box. This is what the computer does with a 2D image in wrapping it around a 3D object.

During the UV unwrapping process, you tell Blender exactly how to map the faces of your object (in this case, a box) to a flat image in the UV/Image Editor window. You have complete freedom in how to do this. (Continuing our previous example, imagine that, having initially laid the box flat on the tabletop, you now cut it into smaller pieces, somehow stretch and/or shrink those pieces, and then arrange them in some way upon a photograph that's also lying on that tabletop ...)

Cartography Example

Cartographers (map makers) have been dealing with this problem for millennia. A cartography (map-making) example is creating a projection map of the whole world. In cartography, we take the surface of the earth (a sphere) and make a flat map that can be folded up into the glove compartment aboard the space shuttle. We 'fill in' spaces toward the poles, or change the outline of the map in any of several ways:



Mercator Projection



Mollweide Projection

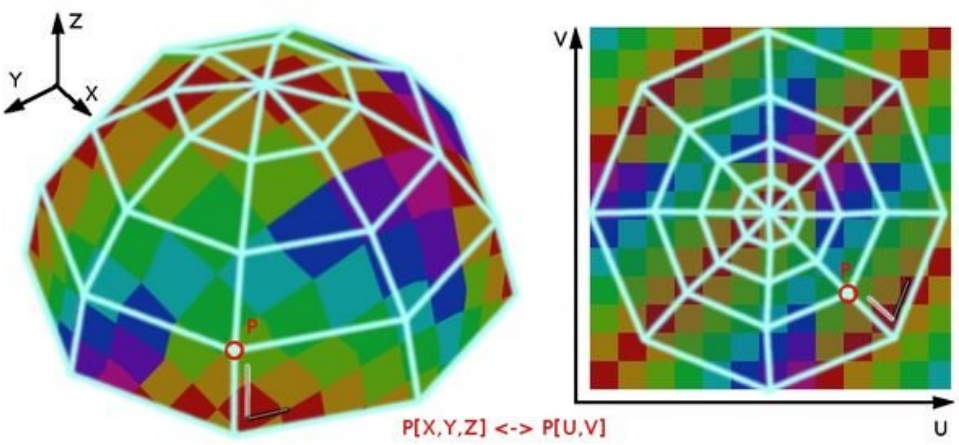


Albers-equal Projection

Each of these is an example of a way to UV map a sphere. Each of the hundred or so commonly accepted projections has its advantages and disadvantages. Blender allows us to do the same thing any way we want to, on the computer.

On more complex models (like seen in the earth map above) there pops up an issue where the faces can't be 'cut', but instead they are stretched in order to make them flat. This helps making easier UV maps, but sometimes adds distortion to the final mapped texture. (Countries and states that are closer to the North or the South Pole look smaller on a flat map than do ones which are close to the Equator.)

Half-Sphere Example



3D Space (XYZ) versus UV Space (click to enlarge)
In this image you can easily see that the shape and size of the marked face in 3D space is different in UV space.

This difference is caused by the 'stretching' (technically called mapping) of the 3D part (XYZ) onto a 2D plane (i.e the UV map).

If a 3D object has a UV map, then, in addition to the 3D-coordinates X, Y, and Z, each point on the object will have corresponding U and V coordinates. (P in the image above is an example of how a point on a 3D object might be mapped onto a 2D image.)

Advantages

While procedural textures (described in the previous chapters) are useful-they never repeat themselves and always "fit" 3D objects-they are not sufficient for more complex or natural objects. For instance, the skin on a human head will never look quite right when procedurally generated. Wrinkles on a human head, or scratches on a car do not occur in random places, but depend on the shape of the model and its usage. Manually-painted images, or images captured from the real world gives more control and realism. For details such as book covers, tapestry, rugs, stains, and detailed props, artists are able to control every pixel on the surface using a UV Texture.

A UV map describes what part of the texture should be attatched to each polygon in the model. Each polygon's vertex gets assigned to 2D coordinates that define which part of the image gets mapped. These 2D coordinates are called UVs (compare this to the XYZ coordinates in 3D). The operation of generating these UV maps is also called "unwrap", since it is as if the mesh were unfolded onto a 2D plane.

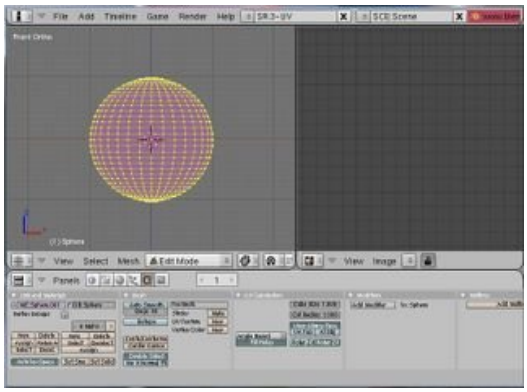
For most simple 3D models, Blender has an automatic set of unwrapping algorithms that you can easily apply. For more complex 3D models, regular Cubic, Cylindrical or Spherical mapping, is usually not sufficient. For even and accurate projection, use seams to guide the UV mapping. This can be used to apply textures to arbitrary and complex shapes, like human heads or animals. Often these textures are painted images, created in applications like the Gimp, Photoshop, or your favorite painting application.

Games
UV mapping is also essential in the Blender game engine, or any other game. It is the de facto standard for applying textures to models; almost any model you find in a game is UV mapped.

The first step is to unwrap your mesh. You want to unwrap when you feel your mesh is complete with respect to the number of faces it needs to have. If you do add faces or subdivide existing faces when a model is already unwrapped, Blender will add those new faces for you. In this fashion, you can use the UV Texture image to guide additional geometry changes.

The UV tools were significantly enhanced for Blender Version 2.46. For users of Blender 2.45 and before, consult this archived page

Getting Started

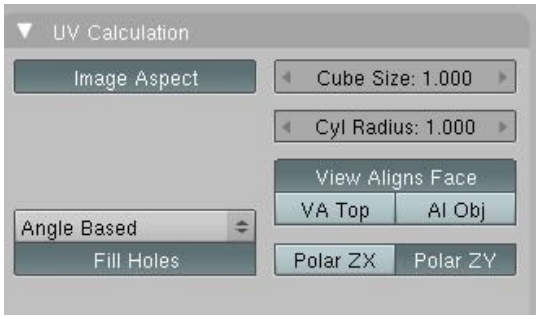


The process of unwrapping your model, called UV Unwrap, is done within Edit Mode in the 3D View window. This process creates one or more UV Islands in the UV/Image Editor window

To begin, choose the SR:3-Material screen layout from the selection list at the top of your screen in the User Preferences window header, and set one of the panes to show you the UV/Image Editor window (Shift+f10), and another pane the 3D window (Shift+f5).

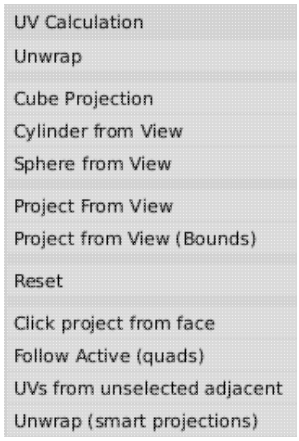
Enter edit mode, as all unwrapping is done in Edit mode. You can be in vertex, face, or edge selection mode.

Notice that the buttons window is in the F9 Editing context.



This panel controls the Unwrap process. For more information on this panel, consult the Reference manual.

Workflow



Choosing the unwrapping method
The process for unwrapping is straightforward, but there are tons of options available, each of which dramatically affect the outcome of the unwrap. By understanding the meaning behind the options, you will become more efficient at unwrapping. The process is:

1. Set the UV Calculation panel to govern the overall unwrap process
2. Select the object to be unwrapped in Object mode
3. Set 'Draw type' to 'Textured' (Alt+z to toggle) in the 3D Viewport window (the 3D Viewport draw type selector is next to the 'Mode' selector in the 3D Viewport's header)
4. Tab into Edit mode, or select Edit Mode from the mode selector on the 3D View header
5. Select the faces you want to unwrap. For ease of use, switch to Face Select mode by clicking the triangle icon in the Select Mode group.
6. In the Mesh panel, activate an existing UV Texture or click the New button next to UV Texture to create a new UV Texture.
7. If you created a New UV Texture, the selected faces are automatically wrapped to *Full Size Reset* (explained later) and are pink if you are in Textured Draw Type mode. Un-selected faces are potato white.
8. Press U or select Mesh->UV Unwrap. The popup menu shown to the right will appear. This menu allows you to choose what algorithm, or method, you want to use.
9. Choose the unwrap method from the popup menu
10. The method will re-calculate the location of unpinned UVs that correspond to vertices for the faces selected. ("Pinning" is also discussed later.)

Every purple dot in the UV map corresponds to a vertex in the mesh. The dotted lines joining the UVs correspond to edges in the mesh. Each face in the UV map corresponds to a mesh face.

Each face of a mesh can have many UV Textures. Each UV Texture will have an individual image assigned to it. When you unwrap a face to a UV Texture in the UV/Image Editor, each face of the mesh is automatically assigned two extra internal features:

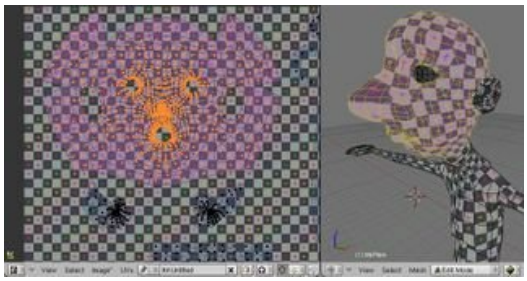
- *four UV coordinates:* These coordinates define the way an image or a texture is mapped onto the face. These are 2D coordinates, which is why they're called UV, to distinguish them from XYZ coordinates. These coordinates can be used for rendering or for realtime OpenGL display.

- *a link to an Image:* Every face in Blender can have a link to a different image. The UV coordinates define how this image is mapped onto the face. This image then can be rendered or displayed in realtime. A 3D window has to be in "Face Select" mode to be able to assign Images or change UV coordinates of the active Mesh Object.

This allows a face to participate in many UV Textures. A face at the hairline of a character might participate in the facial UV Texture, *and* in the scalp/hair UV Texture.

Active Face
The last face clicked is "active" (either selected via right-click or de-selected thru shift-right-click). If you select some faces, then de-select some others, the last face you de-selected is "active". If the active face is a de-selected one, you will not be able to map it or see its UV Texture, since, well, it isn't selected. All other faces that are selected are mapped though; just click on them to see their mapping.

UV Orientation



UV orientation to Face
When you map a face to a UV Texture, you have really no way of knowing the orientation. The flat square face that is a on the XY plane in 3D view could be a trapezoid on its side in the the UV/Image Editor.

The way to compare the orientation of the UV Texture face and the 3D View face is to apply a test image. In the UV/Image Editor window:

1. Sync UV and Mesh Selection by clicking the icon on the UV/Image Editor. A little orange face-link icon shows up in the bottom left corner of the UV/Image Editor window. This shows all UV faces in the UV/Image Editor window, not just the ones selected in 3D View. This gives you a much better perspective on how all the faces are linked and laying flat in the UV/Image Editor
2. Enable View->Draw Faces so that selected faces are pink in the UV/Image Editor which makes them easier to see.
3. Select all the UV faces by A
4. Create a test grid by clicking Image->New. In the popup click UV Test Grid and click OK.

The UV Test grid is a checkerboard with a diagonal series of + in the middle in a specific vertical color rotation: magenta, purple, blue, cyan, 2 x green, yellow, orange, and then back around again. This sequence is offset by one color, so you get a diagonal checker band. This makes it very easy to see the orientation of the UV face relative to the 3D View face.

By RMB selecting the face in either window can you determine the orientation of the window, since the corresponding face in the other window will be highlighted. clicking on a face of the object. If the face was white in the 3D space, and the UV window is black with just a grid background, it means that there is no image mapped to that face. To map an image for a selected face, simply choose it from the image selector or assign a new one. and turns pink when you click, it means that there currently is no UV mapping for it.

UV Rotation

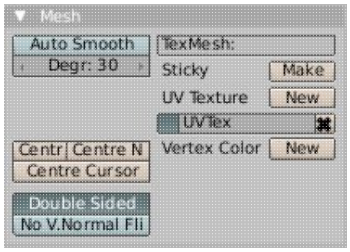
If you don't like the orientation of the UV face, for example if the picture is on its side or upside down, press R to Rotate the mapping of the selected UV's and the corresponding faces. Much like 3D View rotation, the header will show you how many degrees the mapping has rotated. Holding CTRL while rotating constrains the rotation to 5 degree increments.

UV Location

Changing the shape of a UV face (by moving the corner UVs) will also change the orientation of the mapping, since the vertex has not moved, but the mapping of image pixels to the face has (because you moved the UV location). Press G to grab the UV's and move them; LMB to drop, RMB to abort, just like vertices.

You can also snap UV's just like vertices by activating the magnet icon on the UV/Image Editor header. If you do, the Snap Target Mode selector will appear: hold control to snap UVs to the closest, center, or median of another UV island.

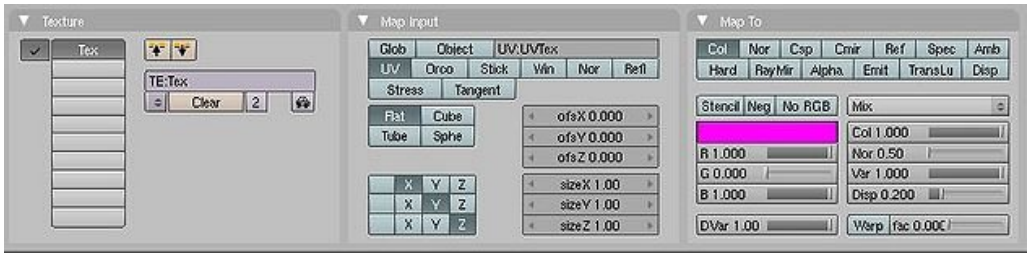
Mapping an Image to Your UVs



When you entered UV Face Select mode, Blender created a mapping in the UV/Image editor window. If you examine your Editing (F9) buttons, the Mesh panel, you will see a new UV Texture entry, called "UVTex" by default. This panel section has a button to the left of the name, used to select the UV Texture map if there are many listed. The name can be changed by clicking into the field and typing. This texture map can be deleted by clicking the X. You can add arbitrary UV Textures by clicking New, but you will rarely have to, since Blender creates them automatically each time you unwrap.

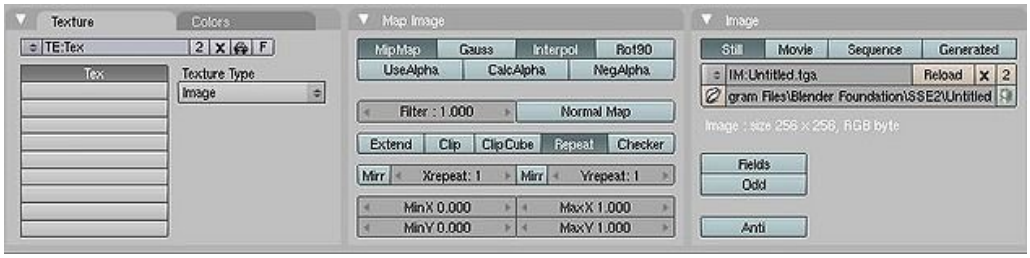
We just very briefly talked about assigning an internally generated image to this UV Layout, the Test Grid. We call that image a "UV Texture" once it is mapped to the mesh using the UV layout and applied as a texture to the material for the mesh. Let's go over that process in more detail.

First, we created the image in memory. With the faces selected in the 3D window, move on over to the UV/Image Editor window. From the header menu, select Image->New. Here you select the size of the image to be used, and/or a test grid. Use the Test Grid to orient yourself to the direction of the UV mapping. If you select Test Grid and click OK, that image of the test grid is mapped to the faces. It cannot be rendered yet, because we have not told Blender how to reflect the light off the object when rendering. We do that by creating a Material Texture.



Material Texture settings

First, Add New material for the object. In the Material Texture tab, click Add New and a texture called "Tex" will be created and assigned to the top channel. In the Map Input panel, specify we want to use our UVTex UV-map, and map the texture to Col (which is selected by default). So, we have told Blender to color the surface of the mesh with a texture, but we haven't told Blender *which* texture to use.



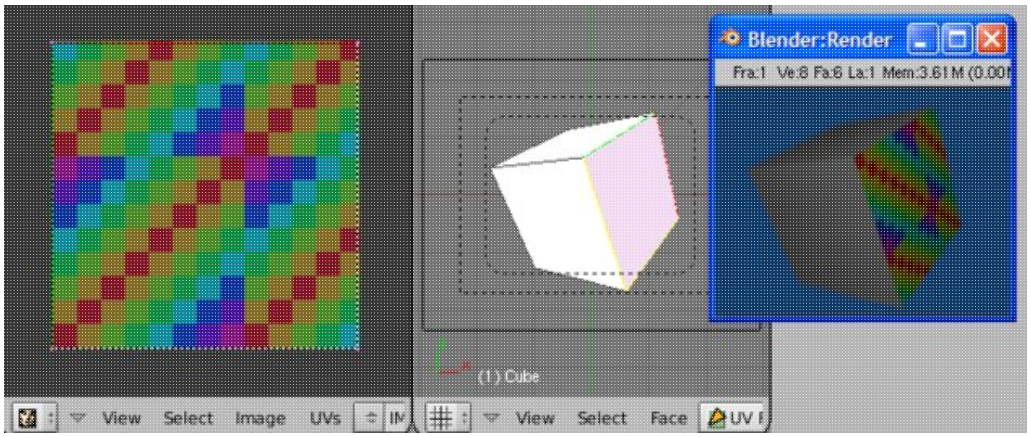
Texture settings
Activate the Texture subcontext buttons, and load the UV image by clicking the the Load button in the Image panel.

A quick render will show that you have automatically mapped the color of the faces of the mesh to an image used as a UV texture. Congratulations!

These three links are crucial to understanding:

1. A UV Layout, shown in the UV/Image Editor, that maps mesh faces to all or part of the image workspace,
2. A Material assigned to those faces which maps a texture to the colors, and
3. An image texture which loads an image.

Now here comes the confusing part: a mesh object can have many UV Layouts (listed in the Editing buttons panel, UV Texture list). A mesh object can have multiple materials, with a single material mapped to one or more faces of the object. A material can have many (up to 10) textures, each assigned to a texture channel. Each of those textures can be an image mapped to a UV Texture. So, Blender provides ultimate flexibility in mapping and layering images to faces of a mesh object.



Note: 3D window in Shaded viewport to show pink selected face.
You do not have to map all faces to an image. If you select the top face of the cube in UV Face Select mode, then Select->Inverse, you will have selected all the faces except the top one. Then, in the UV/Image Editor window, select a different image. Of course, instead of creating a new image, you can load any of the image formats supported by Blender (JPG, PNG, TGA, etc.)

Once a face has been mapped, RMB clicking on that face in the 3D View will cause the UV/Image Editor window to automatically show the image and UV coordinates for the selected face. The UV map for that face will be selected in the Mesh panel.

Images versus Maps

You may have heard the terms "Bump Map" or "Spec Map". These are simply terms to denote how a UV Texture is mapped to the material, and what purpose it serves. A Color image affects the color that you see. For skin, there are four common color maps, or images, that are applied to the mesh through the UV coordinates: Tone, Spec, Blemish, and Makeup. The **Tone** is the basic skin color, **Spec** is the color you see when the light shines directly on the skin. **Blemish** is the wonderful imperfections: varicose veins, pimples, scars, scrapes, and **Makeup** is coloring of the lips, eyes, cheeks. These images are *Map To ... Col,* or color.

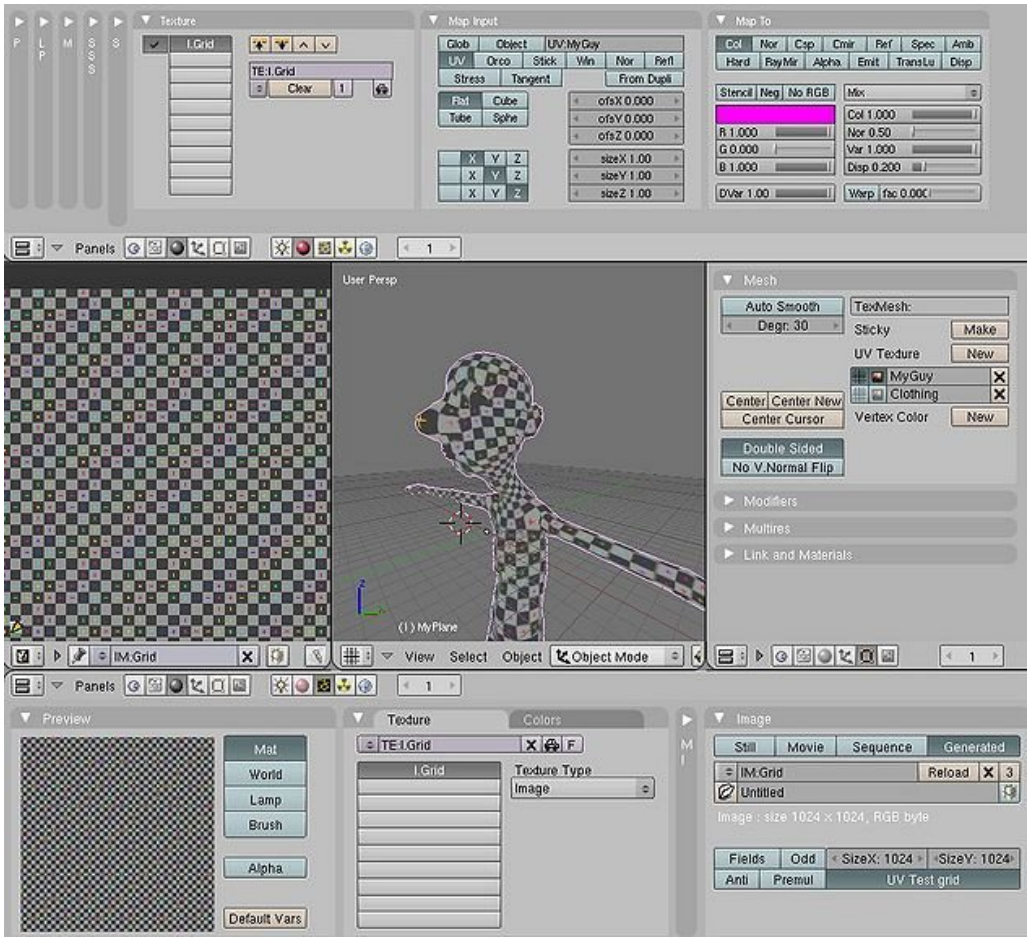
All the other images are gray-scale images, where black (0.0) is "no effect," and white is 1.0 or "full effect." These images map the texture channel indicated, onto the mesh, through the UV Texture. For skin, there are commonly five: Bump (Map to *Nor*), Diffuse (map to *Ref*), Specular (map to *Spec*), Ambient (map to *Amb*), and Hardness (map to *Hard*). The **Bump** map wrinkles and gives the skin pock marks and that cellular unevenness. The **Diffuse** map makes the color weaker in some areas. The **Specular** color indicates the skin's thickness, as thinner skin looks redder because of the blood underneath when exposed directly to light. **Ambient** reflects how flushed the skin is, and **Hardness** highlights where the skin is oily.

For example, the skin is oily around the crease of the nose. For the Hardness map then, the image would be black except where the texture maps to the area around the outline of the nose, and there it would be white. In the Material settings, the maximum hardness might be 50. Instead of being universally hardness of 50 all over, by using the Hardness map we can control where it is between 0 and 50.

Using UV Textures, mapped to some aspect of the texture, enables us to achieve photo realism of any surface through CG. You are not limited to these. I saw a very good zombie skin texture that used an Alpha Map to place holes in the flesh where the bone could show through (ugh).

Panel Interaction

Please excuse me for belaboring the point, but it can be very confusing. There are three different button panel sets going on here that all link to one another, as shown below.



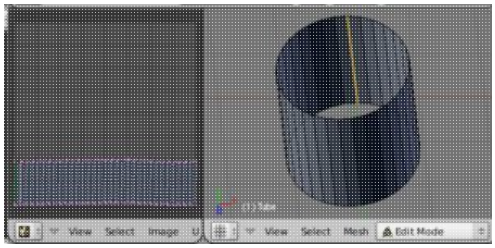
Panel Interactions

In the above example, on top is the Material buttons. The Texture panels for the Material show one texture named I.Grid, mapped to UV using the UV Texture named MyGuy. That I.Grid is mapped to affect the base color of the mesh. On the second row, you see the UV/Image Editor Window showing the Image called Grid, and the window in the center is the 3D View. The next window on the right, middle row, is the Mesh panel from the Editing context buttons, showing that we have two UV Textures; the active one called MyGuy and another one called Clothing. In the bottom row of buttons, we see the Texture subcontext of the Material buttons, and in fact see that the I.Grid texture is an Image texture that loads the Generated image that is 1K x 1K in size.

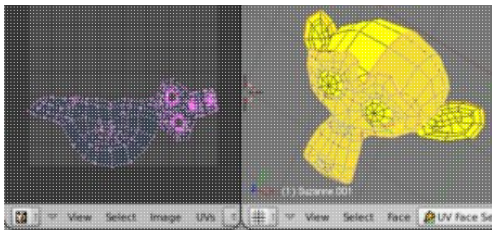
All of these names and such have to match up. I hope you can see how you can use a static image for simplicity, or use an animation to show someone blushing, or beginning to sweat, by using a movie or sequence to affect the different texture channels. I'm sure you've seen the different X-Men effects, so now you know how it was done, and how photo-realism can be achieved.

Unwrapping Using Seams

For most cases, using the Unwrap calculations of Cube, Cylinder, Sphere, or best fit will produce a good UV layout. However, for more complex meshes, especially those with lots of indentations, you may want to define a **seam** to limit and guide any of the unwrapping processes discussed above. Just like in sewing, a seam is where the ends of the image/cloth are sewn together. In unwrapping, the mesh is unwrapped at the seams.



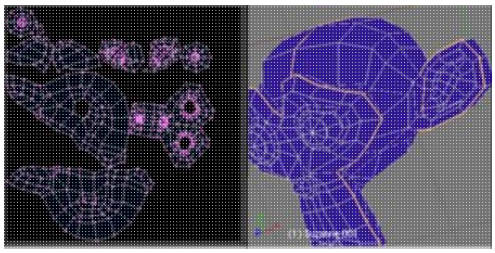
Simple Seam on a Cylinder
The easiest way to define a seam is Edge selection in Edit Mode. Select the edge(s) that define the seam with border or â § ShiftRMB , and press CtrlE or use the Mesh->Edges->Mark Seam menu. In the example to the right, the back-most edge of the cylinder was selected as the seam (to hide the seam), and the default unwrap calculation was used. In the UV/Image Editor window, you can see that all the faces are nicely unwrapped, just as if you cut the seam with a scissors and spread out the fabric.



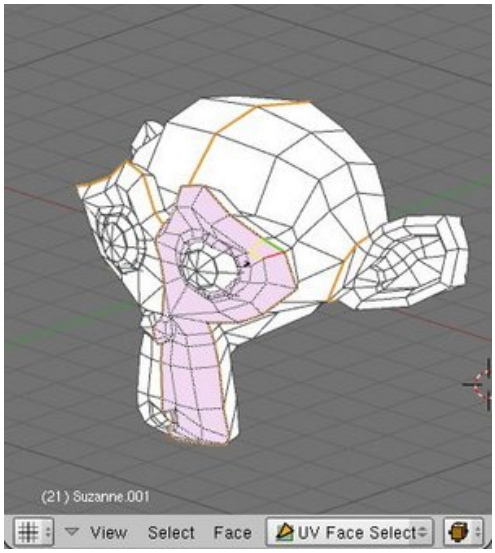
Oops! Forgot an edge in the seam
When marking seams, use the Select->Linked Faces in UV Face Select Mode to check your work. This menu option selects all faces connected to the selected one, up to a seam. If faces outside your intended seam are selected, you know that your seam is not continuous.

To add an edge to a seam, simply select the edge and CtrlE Mark Seam. To take an edge out of a seam, select it, CtrlE and Clear Seam.

You do not have to select continuous seams. If you select one edge, and follow that line and select another edge further down the line using Shift + Alt + RMB â § ShiftAltRMB on the edges in UV face select mode, Blender will connect the line and select all edges in between.



Seamed Suzanne
Just as there are many ways to skin a cat, there are many ways to define seams. In general though, you should think as if you were holding the object in one hand, and a pair of sharp scissors in the other, and you want to cut it apart and spread it on the table with as little tearing as possible. Note that we seamed the outside edges of her ears, to separate the front from the back. Her eyes are disconnected sub-meshes, so they are automatically unwrapped by themselves. A seam runs along the back of her head vertically, so that each side of her head is flattened out.



Face Select Mode.
Another use for seams is to limit the faces unwrapped. For example, when texturing a head, you don't really need to texture the scalp on the top and back of the head since it will be covered in hair. So, define a seam at the hairline. Then, when you select a frontal face, and then select linked faces before unwrapping, the select will only go up to the hairline seam, and the scalp will not be unwrapped.

When unwrapping anything that is bilateral, like a head or a body, seam it along the mirror axis. For example, cleave a head or a whole body right down the middle in front view. When you unwrap, you will be able to overlay both halves onto the same texture space, so that the image pixels for the right hand will be shared with the left; the right side of the face will match the left, etc.

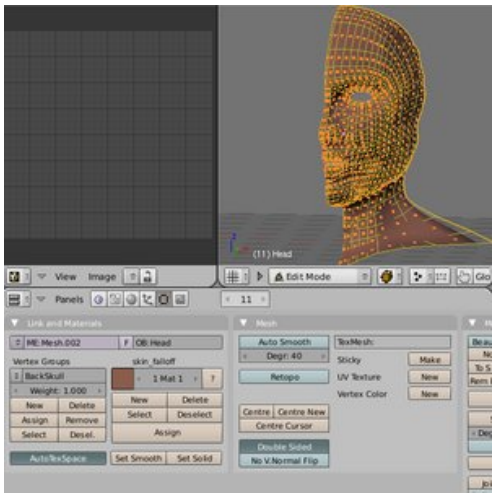
Finally, remember that you *don't* have to come up with "one unwrapping that works perfectly for everything everywhere." As we'll discuss later, you can easily have multiple UV unwrappings, using different approaches in different areas of your mesh.

Unwrapping Multiple Faces



Starting Off in Object Mode
In general, you should only unwrap the faces you need to, and do so in a single unwrap operation. You only need to unwrap faces that will be painted using an image; all other faces can use procedural materials and textures or vertex paint. You want to keep your image as small as possible, so that means you want to keep the

number of faces as small as possible. For example, if the body is always going to be covered in clothes or armor, there is no need to unwrap it. If the back of the head is always going to be covered by hair, there is no need to unwrap the scalp. If you are modeling a chair with an embroidered seat cushion, you only need to unwrap the cushion and not the chair legs. In the example to the right, we only need to unwrap one side of the face, cutting our image size in half, so we leave mirror modifier on; we also do not need to double the number of UV coordinates by applying the Subsurf modifier; we can just leave it as is. Note that at this point, there is no UV Texture in the Mesh panel.

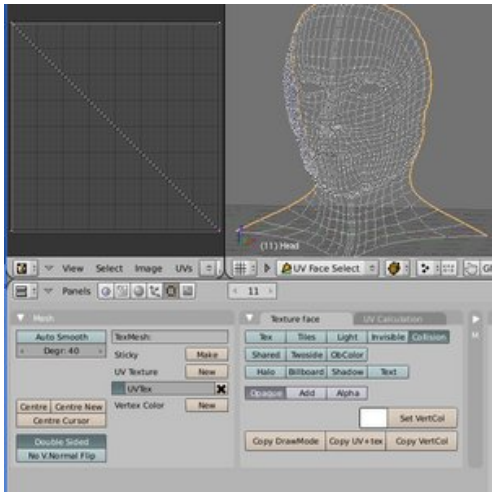


Selecting Faces in Edit Mode

To unwrap multiple faces to a single UV Texture, there are two ways to select the faces you want:

- In Edit Mode (â Tab), enter Face Select mode and select the faces you want.
- In UV Face Select mode using the 3D window, right-click or border select the faces you want.

The example to the right shows that we have hidden many faces from view; the ears and the back of the head. We did so by creating and using the Vertex Groups, selecting the "BackSkull" group and Hiding them from view. We did this because we do not want to unwrap those areas, and we don't want them to get in the way during any further face selection that we may do.



Selecting Faces in UV Face Select Mode

Once we enter UV Face Select mode, Blender creates a UV Texture map for us, shown in the Mesh panel, and displays that map in the UV/Image Editor Window. By default, the map is *full-size reset overlay*, which means that each selected face is mapped to the **full size** of the UV window, and the faces are **overlaid** on top of one another. This situation is identical to the above introduction where you mapped your first face of the

cube; any face of the cube you clicked on would show this same full-size map; any image selected would be displayed on **each** face. If you were to assign an image now, you would see a patchwork style of render, much like the scene in the Matrix when Neo was in the TV room and saw himself tiled over and over. If we were to use an image that was tileable, the surface would be covered in a smooth repetition of that image, with the image skewed to fit the shape of each individual face. Use this unwrapping option to reset the map and undo any unwrapping (go back to the start).

If you did not select all desired faces prior to entering UV Face Select mode, you can select multiple additional faces in any of the following ways:

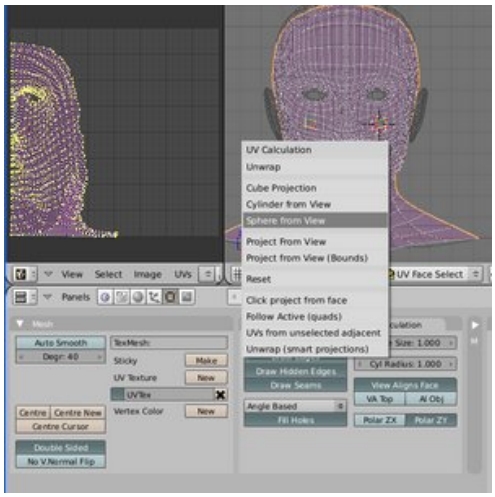
- Press A and all faces of the Mesh will be selected and highlighted by dotted lines.
- Select Linked faces from the 3D View menu; this will select additional faces up to a seam
- You can select many faces by â § ShiftRMB clicking on each one you want. You may spin the object or your view of the object and resume selecting by â § ShiftRMB .
- BorderSelect in the 3D window.
- Enter EditMode and select the vertices or edges that define the face you want. After leaving EditMode and back in UV FAce Select mode, the faces defined by the selected vertices/edges will also be selected.

In UV Face Select mode, just like in Edit mode, you can also hide faces from view so they don't get in the way as you are trying to select faces:

- H to hide selected faces
- â § ShiftH to hide the faces which are *not* selected, sort of the reverse of H by itself.
- AltH to reveal the hidden faces. Note that if you leave and then re-enter UV Face Select mode, they will be unhidden.

Only one face is considered *active*; the last one selected. Since the UV/Image Editor can only show one image at a time, it will show the image for the active face. IF you change the active face mapping (if you move the its UV coordinates), you may be overlapping or affect **all** other face mappings that are part of tha UV map (even if you cannot see them). So, before adjusting a face's UV coordinates, it is best to select adjoining faces to you see the whole map.

Face Unwrap using UV Calculation



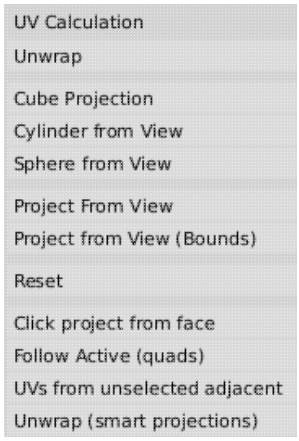
Unwrapping Faces using 3D View Menu

With our faces selected, it is now time to unwrap them to something more useful than Full-Size Reset. We want to unwrap them to the UV/Image Editor window using one of the UV Calculation shown to the right. In

the 3D View header menu, select Face->Unwrap UVs or just press U. You can control the way the faces are mapped in two ways:

- Automatically using the 3D View window command Face->Unwrap UVs U
- Automatically using the UV/Image Editor window command UVs->Unwrap command E
- Creating seams and then unwrapping. A seam is marked in Edit mode by selecting edges that make the seam and then issuing the command to Mark Seam.

Above, when you selected a single face and went into Face Select Mode, Blender automatically mapped the selected face to the entire image for you, based on the face's UV orientation (the red and green borders). You can unwrap a face manually through the Face->Unwrap menu. How the selected faces map over to the image depends on the UV Calculation method that you choose.



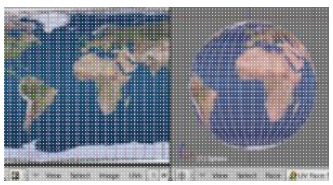
Face Unwrap Menu

The Face->Unwrap->Unwrap option unwraps the faces of the object to provide the 'best fit' scenario based on how the faces are connected and will fit within the image, and takes into account any seams within the selected faces. If possible, each selected face gets its own different area of the image and is not *tucked under* any other faces. If all faces of an object are selected, then each face is mapped to some portion of the image.

This point is crucial to understanding mapping later on: a face's UV image texture only has to use *part* of the image, not the *whole* image. Also, portions of the same image can be shared by multiple faces. A face can be mapped to less and less of the total image.

Based on the fundamental geometry of the object, and how it is viewed, the Face->Unwrap->Cube, Cylinder, and Sphere UV Calculations attempt to unfold the faces for you as an initial best fit. Here, the view from the 3D window is especially important. Also, the settings for cube size or cyclinder radius (Editing buttons, UV Calculation panel) must be set (in blender units) to encompass the object.

Normally, to unwrap a cylinder (tube) as if you slit it lengthwise and folded it flat, Blender wants the view to be vertical, with the tube standing 'up'. Different views will project the tube onto the UV map differently, skewing the image used. Also note the Polar coordinate axis selection buttons in the UV Calculation panel located in the Editing buttons window; they tell Blender which way is up.



Using a Mercator image with a Sphere Unwrap

Recall the opening cartographer's approaching to mapping the world? Well, you can achieve the same here when unwrapping a sphere from different perspectives. Normally, to unwrap a sphere, view the sphere with the poles at the top and bottom. After unwrapping, Blender will give you a mercator projection; the point at the equator facing you will be in the middle of the image. A polar view will give a very different but common projection map. Using a mercator projection map of the earth as the UV image will give a very nice planet mapping onto the sphere.

In the 3D window, Face->Unwrap UVs->Project from View option maps the face as seen through the view of the 3D window it was selected from. It is almost like you had x-ray vision or squashed the mesh flat as a pancake onto the UV map. Use this option if you are using a picture of a real object as a UV Texture for an object that you have modeled. You will get some stretching in areas where the model receeds away from you.

In the 3D window, Face->Unwrap->Reset maps each selected face to the same area of the image, as previously discussed. To map all the faces of an object (a cube for example) to the same image, select all the faces of the cube, and unwrap them using the Reset menu option.

Note: The next four options create a UV map that stretches to maximum, and does not respect the image size. After using these options, you may have to zoom waaaaay out in the UV/Image Editor, select all UV coordinates, and scale them waaaaay down to fit inside the image area.

The Face->Unwrap->Click project from face is a three-step process that lets you orient the projection map according to a selected face's orientation. When selected, the upper User Information bar changes to display the step you are on, temporarily replacing the Blender version and statistics. The first click selects a bottom-left corner of that typical face. The second defines 'up' based on the face orientation, so click on the corner of that face that corresponds to the up direction from the first corner. The third click defines the V coordinate, namely which way is right on the map.

The Face->Unwrap->Follow Active (quads) takes the selected faces and lays them out perfectly square, even if the mesh face is irregularly shaped. Note that it does not respect the image size, so you may have to scale them all down a bit to fit the image area.

Use Face->Unwrap->UVs from unselected adjacent to expand your UV map by selecting un-mapped (new) faces and then this option. Blender will attempt to match up the coordinates of the new face with existing faces previously mapped. This option is helpful if you have added new faces and want to unwrap them.

Finally, the Face->Unwrap->Unwrap (smart projections), (which used to be called the Archimapper) gives you fine control over how automatic seams should be created, based on angular changes in your mesh. This option is especially well suited to unwrapping multiple mechanical objects all at once. The Archimapper is discussed a little later on.

Face Unwrap using the UV/Image Editor Menu

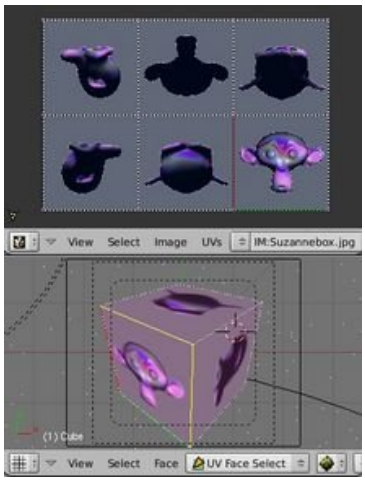
In the UV/Image Editor window menu, the menu item UVs also has a simple Unwrap option. There is no choice of different calculations to use; it is the same as selecting Face->Unwrap UVs->Unwrap in the 3D View. Use this option to unwrap additional faces with an existing set of unwrapped faces. The UV Coordinates for the addtional face will be added to the existing set in a non-overlapping manner. Warning: the UV coordinates for all of the faces will be recalculated.

Texture Face

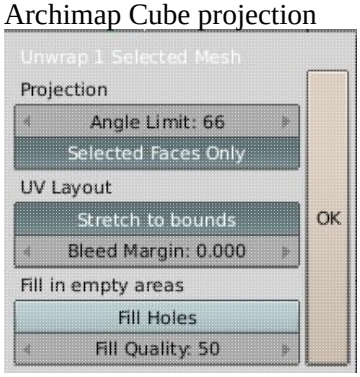


Texture Face Panel
After you have created a UV Map, an additional panel will be displayed in Edit Mode that allows you control over how the image texture is applied to the faces. Consult the Reference manual for more information on this panel.

Unwrap (smart projections)



Archimap Cube projection



Archimap control panel
An architype mapper (archimap) is the final but possibly the best option for unwrapping a mesh, even simple geometric forms. This function examines the shape of your object, the faces selected and their relation to one another, and creates a UV map based on this information and settings that you supply. In prior versions of Blender (pre-2.43), it was a script.

In the example to the right, the Archimapper mapped all of the faces of a cube to a neat arrangement of 3 sides on top, 3 sides on the bottom, for all six sides of the cube to fit squarely, just like the faces of the cube.

For more complex mechanical objects, this script can very quickly and easily create a very logical and straightforward UV layout for you.

Archimap allows you to unwrap multiple meshes in a single step (select multiple meshes in object mode). If you only have one mesh selected and are in Face Select mode, you can unwrap only selected faces, or if you select all the faces of the mesh, the whole object.

The popup panel allows the fine control over how the mesh is unwrapped. The angle limit controls how faces are grouped: a higher limit will lead to many small groups but less distortion, while a lower limit will create less groups at the expense of more distortion.

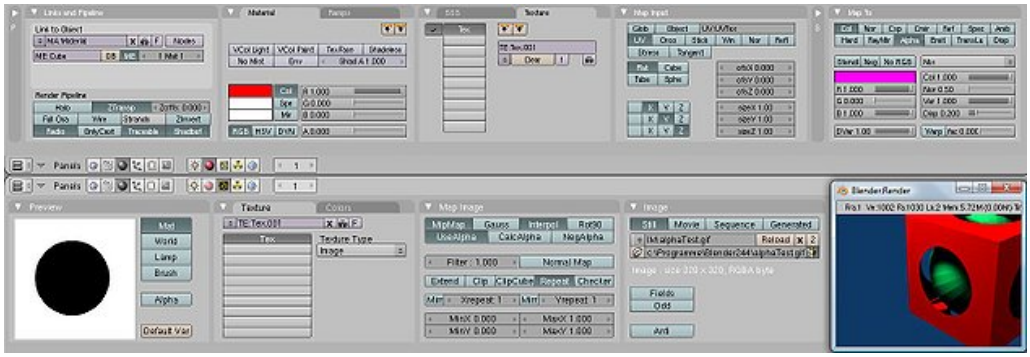
Archimap prefers you to have an image loaded if you want to Stretch to Bounds. When you use an image for a UV Texture, sometimes the edges of the image won't align perfectly or there will may be a border around the outside of the image. The Bleed Margin will shrink the overall UV coordinate map so that this margin is not mapped or used as UV Texture. For complex meshes, Archimap may create (by default) a map that has holes in it, resulting in wasted space in the image. Selecting Fill Holes will cause Archimap to do a little more work at ensuring each face is mapped as large as possible and that there are fewest holes (wasted image texture pixels) possible.

The second step is to work with the UV layouts that you have created through the unwrap process. If you do add faces or subdivide existing faces when a model is already unwrapped, Blender will add those new faces for you. In this fashion, you can use the UV Texture image to guide additional geometry changes.

The UV tools were significantly enhanced for Version 2.46. For users of Blender 2.45 and before, consult this archived page

Activating UV Textures

The UV/Image Editor allows you to map textures directly to the mesh faces. The 3D View window shows you the object being textured. If you set this window into Textured viewport shading, you will immediately see any changes made in the UV/Image Editor window in this window, and vice versa.



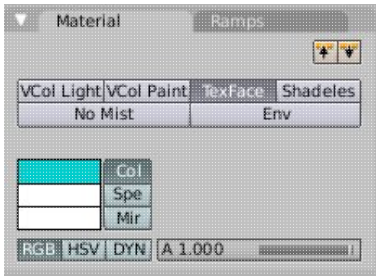
The Material and Texture panels using a UV Texture.
You can edit and load images, and even play a game in the Blender Game Engine with UV textures for characters and object, without a material, and still see them in the 3D window. This is because no 'real' remdering is taking place; it is all just diffuse shading.

To render an image however, you must

1. create a Material for the object, and
2. tell Blender to use the UV Textures on faces when rendering.

To create a Material, you have to click Add New Material in the F5 Shading context.

There are two ways to tell Blender to use the UV Texture when rendering: the Proper way and the Quick Way.



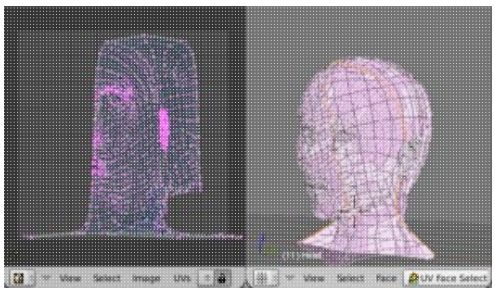
The Material panel with activated TexFace button.

- Quick Way: The quick way is to set up a TexFace Material as shown. To do so, with the buttons window displayed, press F5 to display the Shader Buttons. In the Buttons window, Material settings, click ADD NEW material. On the Material panel, enable TexFace. This way is quick, but bypasses the normal rendering system for fast results, but results which do not respect transparency and proper shading.

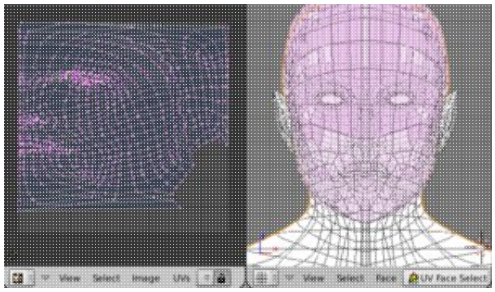
- Proper way: In the Texture channel panel F6((shown above), Add a New Texture which is mapped to the UV texture (it will be mapped to Color by default, and the UV Texture is named "UVTex" by default). Select the Textures subcontext, and define the texture as an image and load the image you want to use. If the image has an alpha channel and you want to use it, click "UseAlpha" in the Map Image panel.

Material is Required for Rendering
You can perform UV Texturing on a mesh within Blender without assigning a material, and you will even see it in your 3D View in textured viewport mode. However, when you render, you will just get a default gray if the object does not have a Material assigned. You will get a black if you do not load an image. If you do not create a texture that uses the image, or enable TexFace, your object will render according to the procedural material settings.

Combining UV Maps

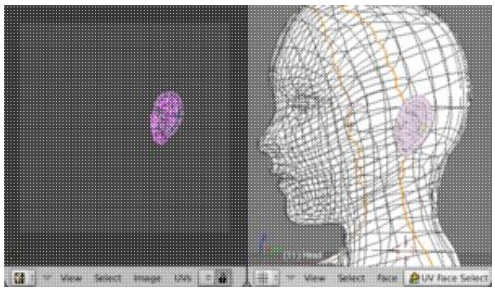


Bad Unwrap-Note Ear and Neck
Very often you will unwrap an object, such as the face example we have been using, and get it 'mostly right' but with parts of the mesh that did not unwrap properly, or are horribly confusing. The picture to the right shows an initial unwrap of the face using the Unwrap from sphere option. The issues are with the ear; it is just a mush of UVs, and the neck, it is stretched and folded under. Too much work to clean up.



Unwrap Face Only, without Ear or Neck

We can tell that the ear would unwrap nicely with just a straightforward projection from the side view, and the neck with a tubular unwrap. So, our general approach will be to unwrap different parts of the object (face, ears, and so on) using different unwrap calculations, selecting each calculation according to whatever works best for that piece. So let's begin: We select only the "face" faces, unwrap them using the *Sphere* calculation, and scale and rotate them somewhat to fit logically within the image area of the UV/Image Editor window pan.

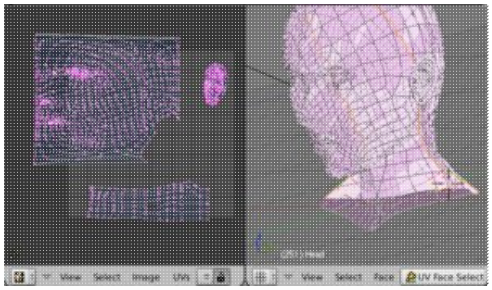


Unwrap Projection: Ear

Once we're satisfied with the face, it's time to turn our attention to the ear. First, unselect the faces you were working with. Their UVs will disappear from the UV/Image Editor, but they are still there, just not shown. (To verify this, you can select a few faces in 3D view and it will show up in the UV/Image Editor.)

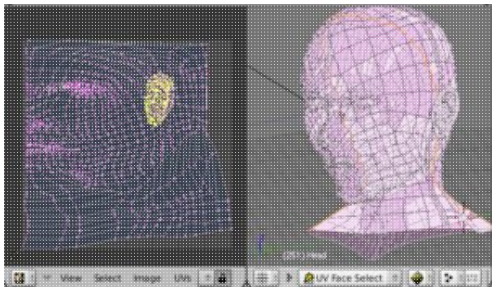
To work on the ear, in the 3D View, we now select only the "ear" faces. You can use Vertex Groups to select the ear faces. Selecting sub-meshes is easy too, since they are not connected to the rest of the mesh. Simply selecting Linked vertices will select that entire submesh. Basically, since you are in edit mode, all of the selecting/unselecting features are available to you.

Now re-unwrap the ear using the *Project* calculation from side view, and scale and rotate them somewhat (discussed in the next section), and place them off to the side. You can do this repetitively, using different UV calculations; each re-calculation just puts those UVs for the selected faces somewhere else. Choose the calculation for each piece that gives you the best fit and most logical layout for subsequent painting of that piece.



UV Maps together

When all of the pieces of the mesh have been unwrapped using the various calculations, you should end up with something that looks like to the Example to the right. All of the sections of the mesh have been mapped, and all those maps are laid out in the same UV Texture map. Congratulations! From here, it is a simple matter of "stitching" (discussed in the next section) to construct the entire UV Map as a single map.



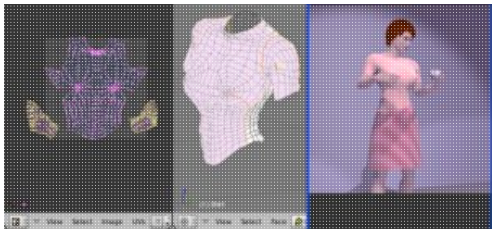
UV Maps Arranged and Stitched

When you have completed arranging and stitching, you will end up with a consolidated UV Map, like that shown to the right, arranged such that a single image will cover, or paint, all of the mesh that needs detailed painting. All of the detailed instructions on how to do this are contained in the next section. The point of this paragraph is to show you the ultimate goal. Note that the mesh shown is Mirrored along the Z axis, so the right side of the face is virtual; it is an exact copy of the right, so only one set of UVs actually exist. (If more realism is desired, the *Mirror* modifier would be applied, resulting in a physical mirror and a complete head. You could then make both side physically different by editing one side and not the other. Unwrapping would produce a full set of UVs (for each side) and painting could thus be different for each side of the face, which is more realistic.)

Averaging UV Islands

The UV Island average tool, (Ctrl+A in the uv window), makes selected UV islands proportionally equal, using their area in the 3d view

Multiple UV Layouts



Mesh with Multiple UV Textures

You are not limited to one UV Layout per mesh. You can have multiple UV layouts for parts of the mesh by creating new UV Textures. The first UV Texture is created for you when you select a face in UV Face Select mode. You can manually create more UV Textures by clicking the New button next to "UV Texture" on the Mesh panel in the Buttons Window, Editing Context) and unwrapping a different part of the mesh. Those faces will then go with that UV Texture, while the previously unwrapped faces will still go with the previous UV Texture. Note that if you unwrap the same face twice or more times (each time to a different UV Texture), the coloring for that face will be the alpha combination of the layers of those UV Textures.

In the example to the right, we have a mesh for a blouse. The mesh has been seamed as a normal blouse would, shown in the middle in UV Face Select mode. Wishing to make a cut pattern, the front of the blouse was unwrapped and basic rotation and scaling was done to center it in the UV/Image Editor window. It was then moved off to the side, while the left and right sleeves were unwrapped, each rotated and scaled. Then, select a sample face from each cloth piece, in the 3D View Select->Linked Faces, and the UV/Image Editor will show all those pieces (as shown to the right). You can then work with all pieces for that UV Texture layout. The example shows all three pieces moved onto the image area for painting. As you can see, the pattern nicely fits a square yard of cloth.

Another UV Layout was created by clicking the New button in the Mesh panel, and the process was repeated for the backs of the sleeves and the back of the blouse. Two images, one for the front and one for the back, are used to color the fabric. In this case, some faces map to the first texture, while other faces map to the second texture.

Additional Options

When you switch your 3D window to UV Face Select mode, Blender changes its menus slightly to give you more and better tools to use.

Similar to Active
: The Select menu now offers a script to help you select items that are similar to the last-selected (active) face. You can pick all the faces, for example, that have the same surface area (plus or minus a limit) as the active one. For each choice, you can add or take away the similar items from the list of faces selected. Criteria include Material, UV Image, Face mode, Vertex colors, and many others. For example, if you were going to change an image used as a UV texture, you would want to select all faces that used that same UV image in order to evaluate the impact of the change on the overall mesh.

Linked Faces
: Recall that not all faces and vertices of a mesh have to be connected. A single mesh can be composed of many disconnected sub-meshes. Selecting Select->Linked Faces is useful for selecting all the faces of a submesh. Also, linking goes up to a seam as previously discussed.

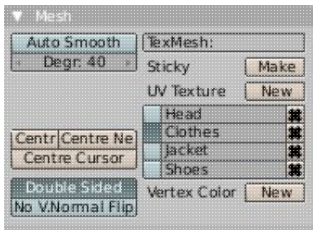
Set Vertex Colors
: If you don't like potato white, select Set Vertex Colors (Face menu) to view your mesh using the active color of Vertex Paint. A blue face color contrasts the seam colors nicely.

Editing Buttons



In the Buttons window, press F9 to get the Editing buttons. When in UV Face Select mode, the buttons change to give you additional options when working with UV maps. The AutoTexSpace button should be enabled.

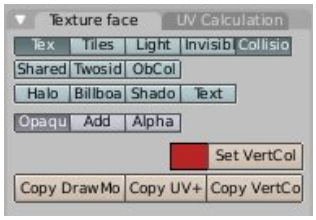
UV Textures List



The Mesh panel (shown to the right) clearly lists the UV Texture maps created for this mesh, and allows you to create New ones as placeholders for future unwrapping operations. Each map has a selector button, a name, and a fat X Delete button. The selected map is displayed in the UV/Image Editor window. The example shows a few UV maps created for a character, and the map for Clothes is selected.

Deleting a UV layout for the mesh destroys all work done in all unwrapping associated the mesh. Click with care. You've been warned.

UV Texture Options



The Texture Face panel is added when you go into UV Face Select mode, so you are not used to seeing it. Lots of neat options for controlling how the UV texture interacts with the rest of your scene. Options for how and if it is shown include: Tex, Tiles, Light, and Invisible. You almost always want it to participate in object and particle Collisions. The texture can be Shared, rendered front and back (Two Sided), and, where blank, can be shown using the object (not the vertex painted) colors.

The texture can color Halos, or be shown like a Billboard, or normally with neither of these selected. The face can used to cast Shadows on other objects.

When rendered, Opaque lets you see this texture, which is almost always a good idea. If you want sections to be transparent where the texture is, use Add; to use the texture as an alpha mask, use Alpha.

Copy UV+ is an action button, not a setting. It copies the UV mapping from the active face (the last face selected) to all other faces that are selected. Select all faces you want similarly mapped first, then â § ShiftRMB the model face, and click Copy UV+



Drawing options (how the UV mapping looks to you while working in Blender) include drawing Faces, Edges, Hidden Edges, and Seams. When you project from a view, normally the View Aligns Face, but you can specify that, for unwrapping round objects (tubes and spheres), the view is aligned to the top (poles) of the object.

By default, Unwrap uses the Angle-Based Formula, or ABF. If you want to use Least-Squares Conformal Method (LSCM) instead, change Angle-Based to Conformal. The reason why ABF is default is because it is

generally better than LSCM. There may be some exceptions for that though, so if an unwrap just isn't clean to you, try conformal.

Please set the cube and tube sizings to encompass your object prior to unwrapping, or incredible skewing and sizing will result.

Saving your UV Layout

The UV coordinates and image links are automatically saved with the mesh in the .blend file. There is nothing extra you need to do.

Saving an outline of your UV Layout

As a way of communicating to an artist who is painting your UV Texture for you, Blender offers a script called Save UV Face Layout (located in the UV/Image Editor Window, UVs->Scripts->Save UV Face Layout) that saves an image in Targa format (.tga) for the object you have selected. The image is an outline of the UV face mapping. The file is (by default) named with the prefix you enter. If Object name is enabled, then the name of the object is appended on so that you don't accidentally overwrite one layout for one object with another's.

Controls allow you to:

Size
: select the size of the image in pixels. The image is always square

Wire
: the thickness of the wire lines that define each face border

All Faces
: if disabled, then only the UV faces selected will be outlined

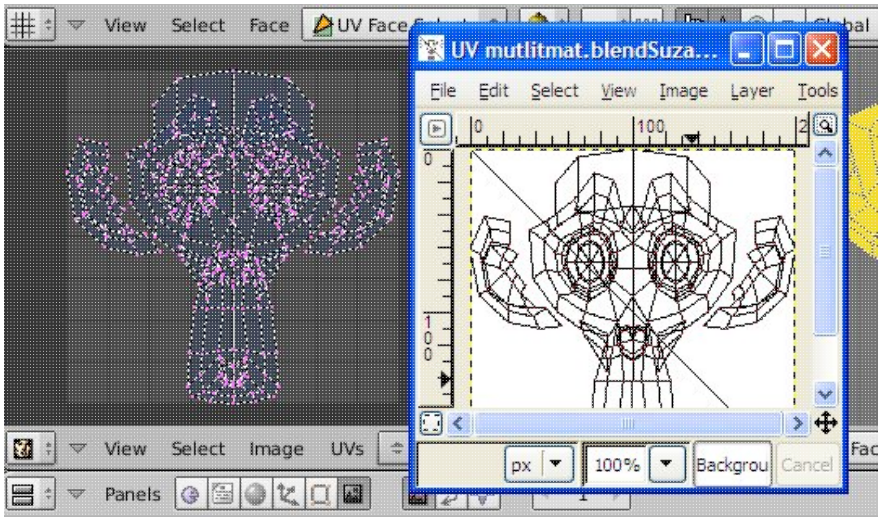
SVG
: if enabled, saves the file as a scalable vector graphic (SVG) format file, not a Targa.

Fill SVG faces
: fills in SVG faces to allow easier editing and face recognition

Edit
: Enabling Edit and specifying the name of an editing program int eh Editor: field will launch that program and load in the saved layout (after the image has been saved).

The image will be lines defining the UV edges that are within the image area of the UV mapping area. Edges outside the boundary, even if selected, will not be shown in the saved graphic.

The artist will use this as a transparent layer in their paint program as a guide when painting your texture. The example below shows Blender in the background, and the Gimp working on the texture, using the saved layout as a guide. Note that targa format supports the Alpha channel, so you can paint transparent areas of the mesh.



Using the layout as a guide in Gimp

After unwrap, you need to arrange the UV maps into something that can be logically painted. Your goals for editing are:

- Arrange the different UV maps together into a coherent layout
- Stitch some pieces (UV maps) back together
- Minimize wasted space in the image
- Enlarge the 'faces' where you want more detail
- Re-size/enlarge the 'faces' that are stretched
- Shrink the 'faces' that are too grainy and have too much detail

With a minimum of dead space, the most pixels can be dedicated to giving the maximum detail and fineness to the UV Texture. A UV face can be as small as a pixel (the little dots that make up an image) or as large as an entire image. You probably want to make some major adjustments first, and then tweak the layout.

Major Adjustments

When you have unwrapped, possibly using seams, your UV layout may be quite disorganized and chaotic. You need to proceed in two steps: Orientation of the UV mapping, and then arranging the UV maps.

3D View: Face Mirror and Rotate UVs

Recall how the red and green outlines show the orientation of the UV Texture relative to the face? Well, you might find that, for example, the image is upside down or laying on its side. If so, use Face->Rotate UVs (in the 3D window in Face Select mode) menu to rotate the UV layout in 90-degree turns. The Face->Mirror UVs to flips the image over like a pankcake in a pan, mirroring the layout and showing you the image 'reversed'.

Merging UV Maps

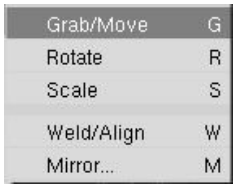
Each time you unwrap, Blender creates a set of UV Coordinates shown in the UV/Image Editor window. So, if you unwrap 2 sides of a cube, and then 3 sides, and then the other remaining side of the cube, there will be 3 sets of UV Coordinates or UV Maps. You would unwrap this way if you wanted to use three different images for each of those maps. However, if you change your mind, simply select all the faces (for example, all 6), and

re-use the Archimapper, the UV/Image Editor window's UVs->Unwrap menu item. The maps will be recomputed to be non-overlapping.

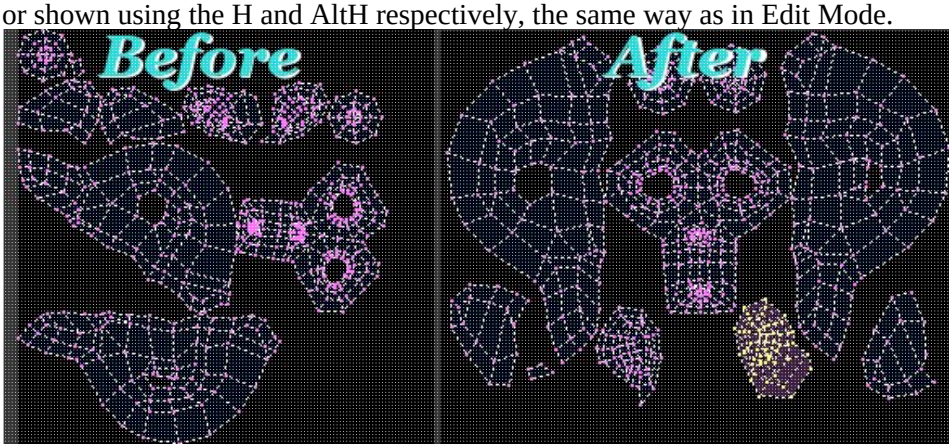
Separating UV Maps

If you want to use a different image for a particular set of faces, but those faces are already mapped to a different UV Texture (image), simply select only those faces and unwrap them again. In the UV/Image Editor window a new map for them will be shown, and you can assign a new image to that map.

Arranging the UV Layout



UV Transformation Menu.
In the UV Editor you will see a representation of your selected faces as yellow or purple vertices connected with dotted lines. Wherever there was a seam or disconnected sub-mesh, there will be a UV map for that piece. During the unwrap, these pieces may be oriented in a haphazard fashion, such that you'll want to rearrange them into a more convenient or sensible arrangement. You can use the same techniques here as in the Mesh EditMode to select, move, rotate, scale, and so on. With the Lock button pressed you will also see real-time feedback in 3D of what you are doing. Scaling and Translating of vertices can be done in the local X or Y axis of the map if needed. Just press X or Y after entering the scale command (S). Proportional editing is also available and it works the same way as in Edit Mode for meshes. Vertices in the UV Editor can be hidden or shown using the H and AltH respectively, the same way as in Edit Mode.



Select Linked, Grab, and Rotate to Arrange layout
When arranging, keep in mind that the entire window is your workspace, but only the UV coordinates within the black square are mapped to the image. So, you can put pieces off to the side while you arrange them. Also, each UV unwrap is its own linked set of coordinates. You can lay them on top of one another, and they will onionskin (the bottom one will show through the top one). To grab only one though, RMB select one of the UV coordinates, and use Select->Linked UVs (CtrlL) to select connected UVs, not border select because UVs from both will be selected.

Deleting a UV Layout

A mesh can have multiple unwraps, each resulting in a UV map for that section of the mesh that was selected. You may get into a situation where you just want to delete everyting and start over. In the Buttons window, the Editing (F9) buttons, there is a Mesh panel. On that panel, click the fat X (Delete) button next to the name

of the UV Texture you want to delete. If you delete every entry there, all the UV layouts for that mesh will be deleted and you will be kicked out of UV Face Select mode. The second that you re-enter UV Face Select mode, you create at least one base UV Texture.

UV/Image Editor Menu

When working with a UV Layout, 99% of your time will be spent in using the UV/Image Editor window. There are many options and features in this window to make you productive. The header consists of a menu, an image selector and a few buttons. Some of the menu items and buttons are relevant when using UV textures. Only the ones relevant to changing the UV layout are discussed in this section.

View Menu

This menu controls how and what you see when working with UVs:

Maximize Window
: Very often, especially when painting, you will want to expand the window pane to full screen to see and work on the details. Or you can wear out your middle mouse button

View All
: Scales the image and UVs to fit within the window pane. Sometimes UVs can stray off into the woods, and this helps you find them and bring them back into the image area.

View Selected
: Centers the view on the selected faces

Update Automatically
: As you move UVs around, the resulting texture changes on the mesh (shown in the 3D window) are updated as you move around. Otherwise, updates are made when you drop the UVs. Use this option if you have the CPU power.

View Navigation
: Hotkeys to zoom in and out the display; same as using the MMB wheel. Also note the window can be panned using â § ShiftMMB . There is no rotate or User view, since we are dealing with a 2 dimensional image.

Draw Shadow Mesh
: Toggles whether to draw an outline of the whole UV Layout for background painting alignment. Very Handy.

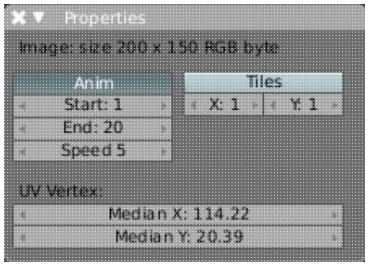
Draw Faces
: When enabled, draws selected faces over the image

Display Normalized Coordinates
: Displays the UV Coordinates in the Properties panel normalized to 1.0. For example, a UV coordiate with X=64 over a 256 grid would be normalized to 0.25.

Composite Preview
: *A work in progress...*

Curves Tool
: Shows the Combined, Red, Green and Blue (CRBG) Color Curves applet. Select a channel (C, R, B, or G) and adjust the curve to affect the colors of the image.

Paint Tool
: Shows the Vertex Paint floating window. See Vertex Paint.



UV Properties

Real-Time Properties
: Shows a floating window which allows you to set Animation and Tile info used in the Game Engine.
- ◊ **Anim** and **Tiles** - See Using UV Textures
- ◊ **UV Vertex** - Normally, the UV Vertex information is in pixels and the black workspace defaults to 256x256; with an image loaded, this display shows you the UV coordinates in relation to the exact image pixel location. With the menu item Display Normalized Coordinates turned on, the vertex location is scaled to 1.0. Origin is the bottom left hand corner of the black workspace. You can click into the X or Y field and manually set the location of a single UV or the median points of multiple UVs, or use the arrows to adjust their location in increments.

Properties
: Many sources can be used for the image portion of the UV Texture. This panel shows you what you are using, and gives you basic controls to select, reload, turn TV interlacing fields on/off (Even or Odd), and do anti-aliasing on the image to smooth the image.

Select Menu

This menu helps you select UVs to work on:

Linked UVs ; CtrlL
: This menu item selects all UVs that are part of the same UV map. Recall that a map is made for every submesh and seamed part of the mesh, and is analogous to a piece of cloth. Selecting Linked UVs works similarly to the command in 3D View. It will select all UVs that are 'connected' to currently selected UVs.

Pinned UVs ; â § ShiftP
: You can pin UVs so they don't move between multiple unwrap operations. This menu item selects them all.

Unlink Selection ; AltL
: Cuts apart the selected UVs from the map. Only those UVs which belong to fully selected faces remain selected following this command. As the name implies, this is particularly useful to unlink faces and move them elsewhere. The hotkey is analogous to the mesh Separate command.

Select/Deselect All ; A
: Selects or de-selects all UV coordinates. When initially unwrapping, you will want to select All UVs to rotate, scale, and move them around.

Border Select Pinned ; â § ShiftB
: Use the box lasso to select only pinned UV coordinates.

Border Select ; B
: Use the box lasso to select normal UV coordinates.

Active Face Select

If turned on, when youRMB  click a 'face' in the UV map, the corresponding face in the 3D View is highlighted. Use this as a nice way of seeing how the initial unwrap corresponds to your real world to aid in rearranging the map into a coherent layout. In addition, (stay with me here) all four UVs that correspond to that face are selected. Very handy for moving/rotating a whole UV face at a time, instead of individual UVs.

Image Menu

I know, I know, you're anxious to get started painting. Refer to the next section for help on this menu.

UVs Menu

This is a long menu that gets cut off sometimes. Here are the UV layout-related choices in the order that they appear:

Scripts
: Displays the list of Python scripts you have loaded on your PC that do handy things with UV Textures:
 - ◊ Auto Image Layout if you have one UV Texture for one part of the mesh, and another different UV Texture (map and image) for another part, use this script to merge them together.
 - ◊ Save UV Face Layout to save a shadow outline of your UV layout; great as a guide layer when using external paint programs. Enable the SVG option to create a file that uses curves.
 - ◊ Texture Baker and UV Painter are discussed in the next painting section.

Show/Hide Faces
: Hide and show selected faces. Works the same way as for meshes in the 3D View, allowing you to focus only on certain parts of the layout.

Proportional Falloff and Proportional Editing
: Same as mesh, and especially useful because often you are moving around a bunch of vertices in a tight space. Hotkey is O. Mousewheel to change the circle of influence, just like mesh editing.

Weld/Align
: When you made a seam and unwrapped it, the mesh was cut at the seam into two pieces. Where they were cut apart, there are now two UV vertices for each mesh vertice; one for a corner or side of one UV map, and another for the other. In order to fill in the gaps between pieces, and make one continuous map, you can Weld them back together. Select the two vertices and then Weld them together. Align vertices in either the X or Y direction so that similar vertices line up. The Weld point is about halfway in between the two.

Mirror
: Occasionally, a UV map will be backwards from your normal way of thinking. This menu option flips the map like a pancake, in either the X or the Y direction.

Transform
: Grab/translate (G), scale (S) and rotate (R) selected vertices. Indespensible, and works just like meshes.

Stitch and Limit Stitch
: Much like the mesh Remove Duplicates function, this function welds together corresponding vertices that are close. Not close as in horseshoes and hand grenades, but close as defined by the Limit. Different parts of a UV map can be stitched if the border UV vertices correspond to the same mesh vertices by using the Stitch command. The Stitch command works joining irregular outlines. Just select the vertices at the border line; if you are using the Stick UVs to Mesh Vertex, then the corresponding coordinates on the other UV Map will also be selected. Stitch (V) snaps together UVs, whatever the distance between them, whereas Limit Stitch (â § ShiftV)only welds UVs within a

given range (Limit:, **20** pixels by default). Its advantage over Weld is that it prevents UVs, that are supposed to stay separate, from being joined together.

Minimize Stretch

Very often, the initial unwrap will have some vertices bunched up right next to some that are spread apart. Much like smoothing out a wrinkled paper, this handy function spreads out bunched up selected vertices. A large face, at an angle to the projection, will result in a small layout. Hence, any image will appear stretched out, like a painting on a balloon. Use this option to relax the stretch a little.

Pin and Unpin

Using the Pin (P) command on selected vertices forces them to stay put between multiple unwrap operations, pinned in their current location. Any subsequent Unwraps will not move them. They appear red and larger than other UV coordinates, like a pushpin. Unpin (AltP) sets them free. Read more about it in the #Using_the_Pin_command section.

Layout Clipped to Image Size

Keeps the UVs in the corral of the Image size. Prevents you from moving a UV outside the image area.

Quads Constrained Regular

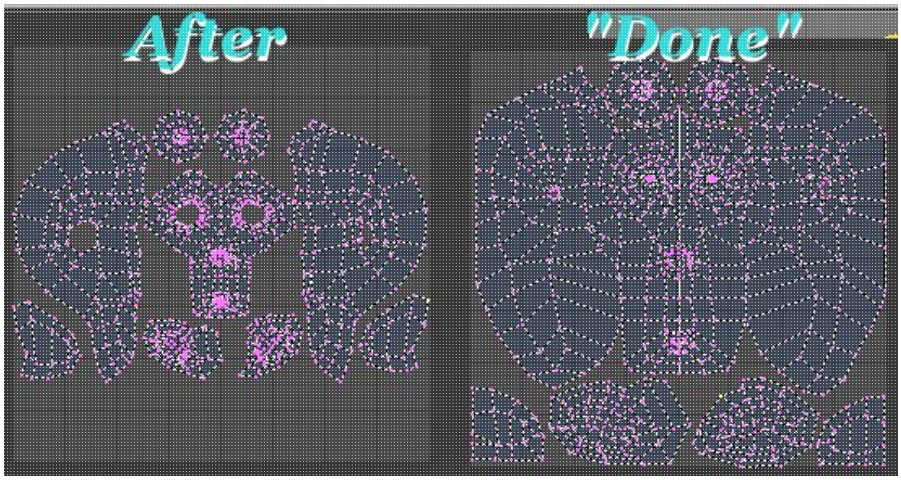
Too geeky for me. Sorry. I mean really, those three words just don't go together in any comprehensible way for my tiny brain :) (In fact, I think this means "Quadrangles (are) constrained (to be) regular"â←¦) When enabled and you grab a UV coordinate, it tries to help you make the UV area square or rectangular. Use it when texturing an orthogonal image to a real-world mesh. Warning: it does an automatic unlink from neighbors if they aren't in line as well.

Snap to Pixels

Put in here by the developers for the perfectionist in all of us; it causes dropped UVs to snap to an even pixel location, so that the UVs line up perfectly with the image. Exactly. Precisely. NO .126 or .436 or any garbage like that. Nice clean units. Gotta love it.

Iteration and Refinement

At least for common people, we just don't "get it right the first time." It takes building on an idea and iterating our creative process until we reach that magical milestone called "Done." In software development, this is called the Spiral Methodology. Applied to Computer Graphics, we cycle between modeling, texturing, animating, and then back to making some modifications to mesh, re-UV mapping, tweaking the animation, adding a bone or two, finding out we need a few more faces, so back to modeling, etc. We continue going round and round like this until we either run out of time, money, or patience, or, in some rare cases, are actually happy with our results.



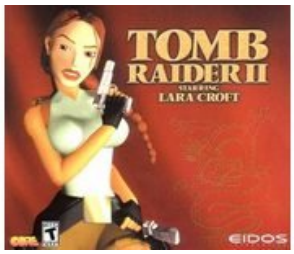
Maximum space for detail areas using Stitch, Align, Minimize Stretch, Scale selected

Now would be a really good time to save an outline of your UV layout. Don't be afraid to try different UV Calculations for the same set of faces to get one that looks good to you and will not be a lot of work to

arrange.

Tradeoffs

For example, consider the creation of a game character. We want the character to look realistic. Realism in appearance for game characters comes in two parts: shape and texture. Regarding shape, we want the model to look proportioned and smooth, and not too blocky. But the more we smooth it out and add features, the more polygons and faces we have. More polys is bad, because it slows down game play because it taxes computer resources. Texture adds realistic colors to the surface, but again, more detailed images means more pixels and thus taxes compute power again. So, there is a balance between the number of faces and the size of the textures, and acceptable realism and game play. Very early games, for example, had a human head the shape of a blocky cube with a texture of a face; let's just say you had to use your imagination. But, it achieved an acceptable level of realism *and* had responsive game play on computers available at the time.



Later on, early 'Tomb Raider' PC games by Eidos are (in my humble opinion) an excellent example of portraying a shapely character in as few polys as possible with small but creative images that faked shadows and creases. Game play was fast, responsive, and a few sharp corners could be overlooked.

Using the Pin command

Developing a game character is thus an iterative creative process. We model the game character, unwrap it, start creating an image for it, and then realize that we have too many faces. We reduce the faces, unwrap again, and start repainting.

This is where the **Pin** command comes in. The first time we unwrap and start drawing an image (for example the clothes), we want to preserve that work. When we unwrap the second time, we don't want the chest UVs to be mapped somewhere else on the image; we want them to stay right where we put them, right over the painting of the vest. So, before unwrapping the second time, we Pin them, then change the mesh, and unwrap the mesh again. The second time we unwrap, those pinned UVs coordinates will stay in place, no matter what UV Calculation method we use or how we change seams.

The other use of the Pin command is sort of an undo. Imagine that when editing the UV map, you discover that yesterday you accidentally welded five unrelated UV coordinates together by mistake. Undo is therefore not available, and you don't want to have to go back to an old copy. Don't feel stupid; your author did so in generating these examples. So you pin all the coordinates except the welded one, re-unwrap, and the UV coordinates will be 'restored'.

Refining the Layout

Refinement comes into play when we finally look at our character, and realize that we need more detail in a particular spot. For example, areas around the eyes might need crow's feet, or we need to add a logo to the vest. As you start to edit the image, you realize that there just aren't enough pixels available to paint the detail that you want. Your only choice is to expand the size (scale out) that UV face. Using the minimize stretch or scale commands, you expand the UV faces around the eyes or chest, allocating more pixels to those areas, but at the same time taking away pixels (detail) from something else, like the back of the head. After refining the UV map, you then edit the image so that it looks right and contains the details you want.

Reusing Textures

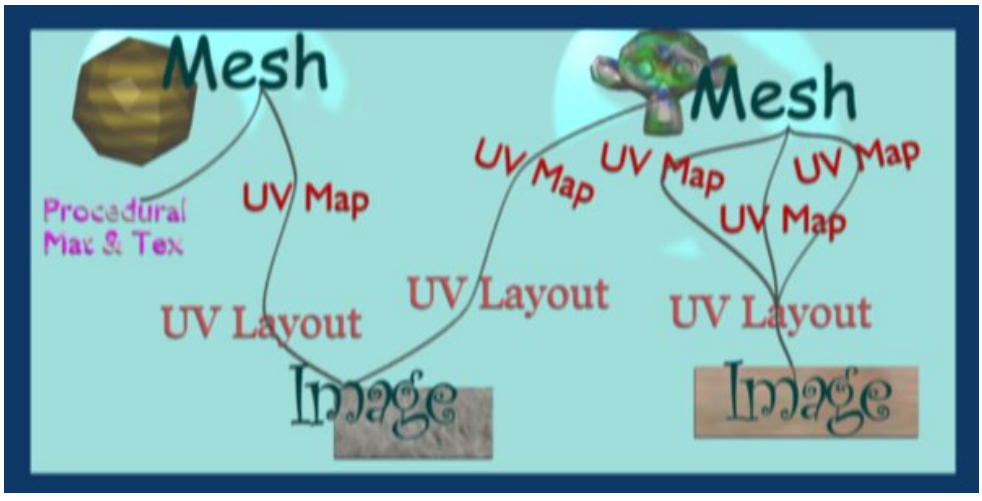
Another consideration is the need to conserve resources. Each image file is loaded in memory. If you can re-use the same image on different meshes, it saves memory. So, for example, you might want to have a generic 'face' painting, and use that on different characters, but alter the UV map and shape and props (sunglasses) to differentiate. You might want to have a "faded blue jeans" texture, and unwrap just the legs of characters to use that image. It would be good to have a generic skin image, and use that for character's hands, feet, arms, legs, and neck. When modeling a fantasy sword, a small image for a piece of the sword blade would suffice, and you would Reset Unwrap the sword faces to re-use that image down the length of the blade.

Up to this point, we have only talked about half of UV Textures; the UV map and layout. The layout is the arrangement of all the UV maps. Each UV map 'maps' image pixels to a mesh face. There is one UV map for each seam or sub-mesh. The entire layout is colored by an image. Blender provides several features that help you when working with the image part of the UV Texture. Blender also includes a built-in texture painting program. This section discusses how to use images effectively.

Goals

Texture images take up precious memory space, often being loaded into a special video memory bank that is very fast and very expensive, so it is often very small. So, keep the images as small as possible. A 64x64 image takes up only one fourth the memory of a 128x128 image. For photo-realistic rendering of objects in animations, often larger image textures are used, because the object might be zoomed in on in camera moves. In general, you want to use a texture sized proportionally to the number of pixels that it will occupy in the final render. Ultimately, you only have a certain amount of physical RAM to hold an image texture and the model and to provide work space when rendering your image.

If you can re-use images across different meshes, this greatly reduces memory requirements. You can re-use images if you map those areas of the meshes that "look alike" to a layout that uses the common image. In the overview below, the left image is re-used for both the sphere and a portion of the monkey. The monkey uses two layouts, one which has one UV map of a few faces, and another that has three maps.



How all the parts of UV Texturing work together
You don't have to UV map the *entire* mesh. The sphere above on the left has some faces mapped, but other faces use procedural materials and textures. Only use UV Textures for those portions of your mesh where you want very graphic, precise detail. For example, a model of a vase only needs UV Texture for the rim where decorative artwork is incorporated. A throw pillow does not need a different image for the back as the front; in fact many throw pillows have a fabric (procedural material) back.

As another example, you should UV map both eyes of a head to the same image (unless you want one bloodshot and the other clear). Mapping both sides of a face to the same image might not be advisable, because the location of freckles and skin defects are not symmetrical. You could of course change the UV map for one side of the face to slightly offset, but it might be noticeable. Ears are another example where images or section of an images can be mapped to similar faces.

Workflow

The process is to:

1. Create the Mesh. Unwrap it into one or more UV Layouts
2. Create one or more Materials for the Mesh
3. Create one or more images for each UV Layout and aspect of the texture
 1. Paint directly on the mesh using Texture Paint in the 3D window
 2. Load and/or edit an image in the UV Editor window
 3. Bake the existing materials into an image for the UV Editor window
4. Apply those images as UV Textures to the mesh to affect one or more aspects of the mesh
 1. Map to Color to affect the diffuse coloring of the mesh
 2. Map to Nor to give the surface a bumpy or creased look
 3. Map to Spec to make certain areas look shiny and oily
 4. Many other Map To options are available for photo-realistic results
5. Layer the UV Textures to create a convincing result.

Using Images and Materials

For an image to use as the color and alpha (transparency) of the texture, you can create an image in an external paint program and tell the UV/Image Editor to Open that file as the texture, or create a New image and save it as the texture.

If you want to start off by creating an image using an external paint program, you will want to save an outline of your UV faces by using the script Save UV Face Layout located in the UVs menu. This script was discussed here.

Creating an Image Texture

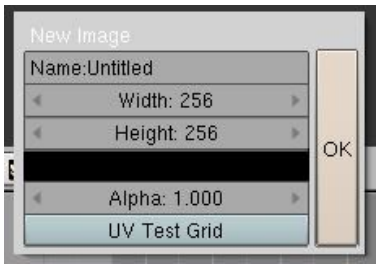
To create an image within Blender, you have to first create a New Blank Image with a uniform color or test grid. After that, you can color the image using the:

- Vertex colors as the basis for an image
- Render Bake image based on how the mesh looks in the scene
- Texture Baker script (discontinued past version 2.43)

After you have created your image, you can modify it using Blender's built-in Texture Paint or any external image painting program.

See Texture in 3D View but does not Render
You may be able to see the texture in Textured display mode in the 3D View; this is all that is required to have textures show up in Blender's Game Engine. Rendering, however, requires a material. You must have a TexFace material assigned to the mesh for it to render using the UV Texture. In the Material settings, ADD NEW material to a selected object and enable TexFace.

Creating a New Blank Image



Once you have unwrapped, use the UV/Image Editor Image->New to create a new image that will be the UV Texture.

Name
: Give your image a more descriptive name

Width and Height
: If you are texturing an item for a game, it is best to use image sizes that are powers of two (16, 32, 64, 128 ...) for both width and height, so they can be drawn properly in real-time using OpenGL. Most 3D cards don't support images larger than 2048x2048 pixels.
For rendered artwork, textures can be any size. Images do not have to be square; they can be any size that you want, provided you have the memory and graphics display capability. Size should be based on the amount of detail that you need.

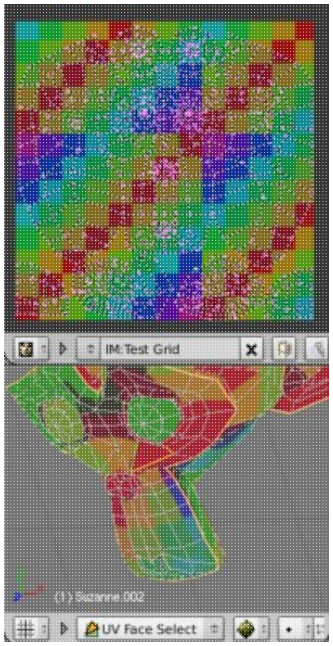
Base Color
: Click on the color swatch (default is black) to bring up the color picker applet, and choose the starting base color.

Alpha
: Choosing an alpha less than 1 will allow other layers of UV textures to show through. The key to realistic looking skin, for example, is many layers that can be individually controlled.

UV Test Grid
: Click this to use the UV Test grid discussed below.

Of course, this simplest of new images will be blank. You will have to start texture painting to get some color into those cheeks. Esc aborts the new image creation.

Using the Test Grid



Use the UV Test Grid option to check for undue stretching or distortion of faces. If your image is a base uniform pattern and you want the application of that image to your model to look like cloth, you do NOT want any stretching (unless you want the cloth to look like spandex).

When you render, the mesh will have the test grid as its colors, and the UV Texture will be the size image you specified. You can save the UV image using the Image->Save menu.

Modifying your Image Texture

To modify your new Texture, you can:

- UV Painter script to create an image from vertex colors
- Render Bake an image based on how the mesh looks
- Texture Baker script (version 2.43-) to create an image
- Paint using Texture Paint
- Use an outside paint package to create an image
- Use the "projection painting" feature of recent versions of Blender
- Some combination of the above.

The first three options, (UV Painter, Render Bake, and Texture Baker) replace the image with an image that they create. Texture paint and outside packages merely add or enhance the image. Regardless of which method you use, ultimately you must:

- Save your texture in a separate image file (for example JPG for colors, PNG with RGBA for alpha)
- *(optionally)* Pack the image inside the blend file (UV/Image Editor Image->Pack as PNG)

The advantage to saving as a separate file is that you can easily switch textures just by copying other image files over it, and you can use external editing programs to work on it. The advantage to packing is that your whole project is kept in the .blend file, and that you only have to manage one file.

Creating an image from Vertex Colors via UV Painter Script

Removed in Blender ver 2.43, TODO: Equivalent feature?

The UV Painter script is located in the UV/Image Editor UVs menu. This nifty script takes the vertex colors from the selected faces and colors in your UV layout. The line button draws an outline of your UV map, and the Scale button scales up or down the image. When you Save your image, a targa-format image file is created for your editing pleasure. UV Painter is an easy way to start painting your mesh, once you have finished your layout.

The Redraw button updates the painting based on any UV layout changes or vertex painting done since the last time the cursor visited the window, although updates may occur automatically.

A small (bug?) Note: After saving the targa image, you must edit it with an external program (Gimp) as it saves dead space off to the side, and you want to scale it in pixels.

Replacing an image via Render Bake

The Render Bake feature provides several tools to replace the current image based on a render of:

- Vertex Paint colors
- Normals (bumps)
- Procedural materials, textures and lighting
- Ambient Occlusion

Click on the hyperlink above for more info. The Render Bake produces an image, already mapped to your UV layout, shown in the UV/Image Editor.

Creating an image via Texture Baker Script

Removed in Blender ver 2.43, TODO: Equivalent feature?

Blender 2.43+
The Texture Baker script, available from the UV/Image Editor UVs menu, saves a UV texture layout of the chosen mesh, that can be used as a UV map for it. It is a way to export procedural textures from Blender as normal image textures that can be edited with a 2d graphical image manipulation program or used with the mesh in games and other 3d applications.

Select your faces in the 3D View Face Select mode; run the script by selecting the menu item. A pop-up panel will appear, allowing you to change the image name or leave it be. The first time you run the script, supply the filename. Choose a reasonable resolution, and an image is rendered. The image depends on your Material settings:

- If you do not have either VCol Paint or TexFace enabled in your material buttons, you will get the UV Layout colored with the procedural (base) material and textures that are current; for example a shadeless purple that is marbled.
- With VCol Paint enabled, the rendered image includes the procedural materials and textures, modulated by the vertex paint job.
- Selecting both VCol Paint and TexFace incorporates procedurals, vertex, and current UV image, melding them all into one beautiful render.

Save the image using Blender's File->Save Image dialog, and the image will be saved in the format specified in the render settings. You can then load it as discussed below.

Broken in 2.44
In 2.43 and before, if the script does not run completely perfectly with you answering all questions, it may leave all Layers unselected, with your 3D View thus blank. Don't know why, it just does.

Create/Modify via an External Image Paint Program

Using your favorite image painting program, draw something that matches your UV layout. Then save your changes, and back in Blender, use the Image->Open menu command to load it as your UV image for the mesh in Face Select Mode for the desired (and active) UV Texture layer.

Modify an image via Texture Paint

Use the UV/Image Editor menu *Image->New*. Then start painting your mesh with Texture Paint.

Using a Saved Image

Use Image->Open to search for and use images in any popular format as the UV Texture. It will be opened and placed as a background to the layout. More often than not, use an image that follows your UV layout. Recall that you saved an outline of your layout and used that to guide the painting of your UV texture image.

If you open an avi file, the first frame of the animation will be used. You can not open an image sequence or use a frame in the middle of a file.

Finding Images

When you open the .blend file, Blender goes out and loads in the most recent image from its location on your hard drive.

In a team environment, or if you are using an external paint program to edit the image while the .blend file is active, and the file is updated and re-saved, use the UV/Image Editor to Image->Reload it and see the latest and greatest in Blender. Also, use Reload if you have mapped more faces to an image, and the 3D View will be updated with the latest image mapping back to faces.

If you move the image file, Blender may not be able to find it, and you will have to Image->Replace it. Use this option to map a UV layout to a different image altogether.

Organize your images
In any scale project, there quickly becomes hundreds of images that are used as UV Textures. Set up a //tex/UV/ directory to hold your UV Textures. Practice good directory and change management. Unfortunately, Blender does not support the Windows shortcut links.

Blender contains a Find Image Target Paths script in the UVs menu in the UV/Image Editor window. Starting from a specified *root* directory for your project, this script will search down and, based on filename, reload images automatically. Use this script if you have renamed a subdirectory or moved a few images around inside your project.

The other way to locate images is to change one of your Blender UI panes to an Image Browser window. This file browser shows you thumbnails and image information (size, format, etc.) of only image files in a directory.

Mapping the Image Texture

Some Modifiers Prevent UV Mapping
In particular, the Decimate Modifier, even if it is only in the Editing modifier list (stack) and not actually applied to the mesh, prevent UV mapping, since it affects the number of vertices and thus UV coordinates.

You then have to create a new Material for the mesh. Then you have two ways of applying that texture to the material:

- The proper way is to map the image using the UV texture, and to load that image as an image texture.
- The quick way is to enable *TexFace* in the Material panel. This tells Blender to use the UVTexture as the base material color (alpha is ignored). Any other textures are layered on top of that base from the top texture channel to the bottom according to their mix method. For example, a wood texture mapped to Alpha on top of a TexFace image makes streaks of the material transparent.

To map the image as a Texture channel (usually the top channel), in the Map Input sub-panel enable UV, and enter the name of the UV Texture (by default UVTex). The advantage is many more options in the Texture's Map Image panel, such as UseAlpha and repeat/mirror. You can also control how multiple textures are layered onto one another by their order in the Texture channels, as well as how they mix with each other, animate their color influence, etc.

The previous pages explained how to create a set of UV Layouts for different portions of the mesh. For example, there may be one UV Layout for the face of a character, and another for their clothes. Now, to texture the clothes, you need to create an image at least for the Color of the clothes, and possible a "bump" texture to give the fabric the appearance of some weave by creating a different image for the Normal of the clothes. Where the fabric is worn, for example at the elbows and knees, the sheen, or Specularity, of the fabric will vary and you will want a different image that tells Blender how to vary the Specularity. Where the fabric is folded over or creased, you want another image that maps Displacement to the mesh to physically deform the mesh. Each of these are examples of applying an image as a texture to the mesh.

As another example, the face is the subject of many questions and tutorials. In general, you will want to create a Material that has the basic skin color, appropriate shaders, and sub-surface scattering. Then you will want to layer on additional UV Textures for:

- Freckle map for Color and Normal aspects
- Subdermal veins and tendons for Displacement
- Creases and Wrinkles and skin cell stratification for Normal
- Makeup images for Color
- Oily maps for Specularity
- For a zombie, Alpha transparency where the flesh has rotted away *(ewwww....)*
- Under chin and inside nostrils that receive less Ambient light
- Thin skin is more translucent, so a map is needed for that

Each image is mapped by using another Texture Channel. Each of these maps are images which are applied to the different aspects (Color, Normal, Specularity) of the image. Tileable images can be repeated to give a smaller, denser pattern by using the Texture controls for repeat or size.

Replacing the active Image

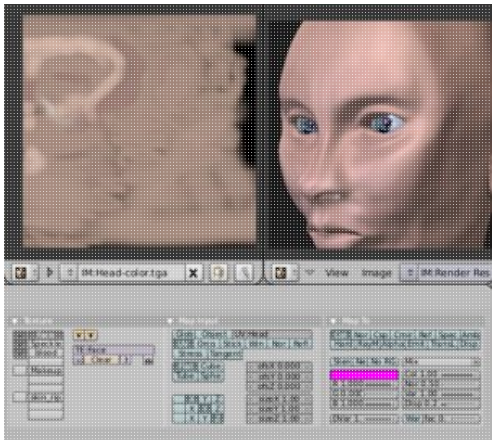
Recall that each face gets coordinates and a link to an image. To map a face to a different image, simply select that face (or faces) and use the UV/Image Editor window Image menu to Replace the current image with an existing file (such as a JPG or PNG file).

Packing Images inside the Blend file

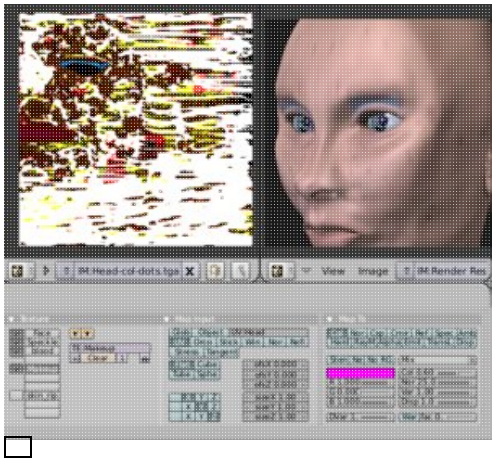
If you pack your .blend file, the current version of all UV Texture images are packed into the file. If those files later change, the updates will not be automatically re-packed; the old version of the image is what will be used. To update, you will have to re-pack or reload.

The File->Append function automatically goes into .blend files and shows you the image textures packed in it. The public domain Blender Texture CD is also a great resource, and there are many other sources of public domain (and licensed) textures. All textures on the Elephant's Dream CD are public domain. And if it looks like a duck and quacks like a duck...

Layering UV Textures



Base UV Texture



Layered UV Texture

Great textures are formed by layering images on top of one another. You start with a base layer, which is the base paint. Each successive layer on top of that is somewhat transparent to let the bottom layers show through, but opaque where you want to add on to details.

To avoid massive confusion, all image textures for a mesh usually use the same UV map. If you do, each image will line up with the one below it, and they will layer on top of one another like the examples shown to the right. To do this, just create one UV Texture (map) as described in this section. Then, create material image textures as described in the procedural materials section. Instead of mapping to Original Coordinates (OrCo), map to UV. Use that map name repeatedly in the Material->Textures->Map Input panel by selecting UV and typing the name in the text field. In the example to the right, our UV Texture is called "Head" (you may have to expand the image to see the panel settings). Then, the image texture shown will be mapped using the UV coordinates. In the "Base UV Texture" example to the right, the face has two textures UV mapped; one for a base color, and another for spots, blemishes and makeup. Both textures use the same UV Texture map as their Map Input, and both affect Color. The Makeup texture is transparent except where there is color, so that the base color texture shows through. Note that the colors were too strong on the image, so they amount of Col affects is turned down to 60% in the second layer (the blemish layer).

Normally, we think of image textures affecting the color of a mesh. Realism and photo-realistic rendering is a combination of many different ways that light interacts with the surface of the mesh. The image texture can be Mapped To not only color, but also Normal (bumpiness) or Reflection or any of the other attributes specified in the Map To panel. If you paint a grey-scale image (laid out according to the UV Layout) with white where the skin is oily and shiny, and dark where it is not, you would map that input image according to the UV Layout, but have it affect Specularity (not color). To make portions of a mesh transparent and thus reveal another mesh surface underneath, you would paint a grey-scale image with black where you want the texture transparent, map input to UV, and map it to Alpha (not color). To make portions of a mesh, like a piece of hot metal, appear to glow, you would use a grey-scale image mapped to Emit.

Believe it or not, this is only "the tip of the iceberg!" If everything that's been described here just isn't enough for you, the *texture nodes* feature, introduced in recent versions of Blender, enables you to layer and combine textures in almost any way you can imagine.

Mix and Match Materials



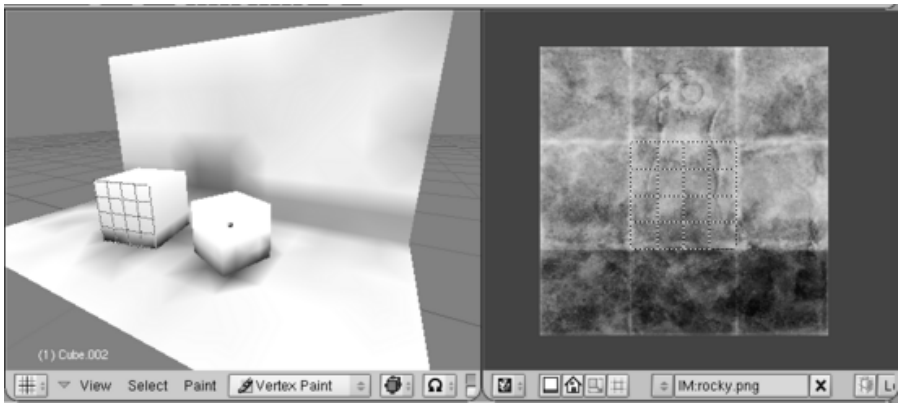
You can mix and match procedural materials and textures, vertex paint, and UV textures onto the same mesh.

The image to the right has a world with a red ambient light. The material has both VCol Paint and TexFace enabled, and receives half of ambient light. A weak cloud texture affects color, mixing in a tan color. The right vertices are vertex painted yellow and the left is unpainted procedural gray. The UV Texture is a stock arrow image from the public domain texture CD. Scene lighting is a white light off to the right. From this information and the User Manual thus far, you should now be able to recreate this image. In other words, I have taught you all that I know, my young padawan. May the Force be with you. Oh wait, there's more...

You can also assign multiple materials to the mesh based on which faces you want to be procedural and which you want to be texture-mapped. Just don't UV map the faces you want to be procedural.

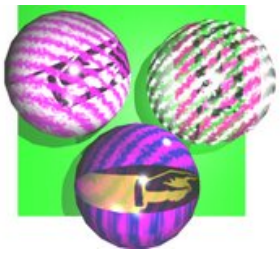
You can use UV Textures and VertexPaint (V in the 3D View window) simultaneously, if both are enabled in the Material settings. The vertex colors are used to modulate the brightness or color of the UV image texture:

- UV Texture is at the base *(TexFace)*
- Vertex paint affects its colors, then
- Procedural textures are laid on top of that,
- Area lights shine on the surface, casting shadows and what not, and finally
- Ambient light lights it up.



Vertex colors modulate texture.
A UV Layout can only have one image, although you can tile and animate the image. Since a layout is a bunch of arranged UV Maps, and a UV Map maps many mesh faces, a face can therefore only have one UV Texture image, and the UV coordinates for that face must fit entirely on the image. If you want a face to have multiple images, split the face into parts, and assign each part its own image. *(Or* you can get fancy with Nodes, but that's another story ...)

Using Alpha Transparency



Alpha UV Textures
Alpha 0.0 (transparent) areas of a UV Image render as black. Unlike a procedural texture, they do not make the base material transparent, since UV Textures do not operate on the base procedural material. The UV texture overrides any procedural color underneath. Procedural Textures are applied on top of UV Textures, so a procedural image texture would override any UV Texture. Transparent (black) areas of a procedural texture mapped to alpha operate on top of anything else, making the object transparent in those places. The only thing that modulates visible parts of a UV Texture are the Vertex Colors. In the example to the right, the finger image is transparent at the cuff and top of the finger and is used as a UV Texture. All three balls have a base material of blue and a marbling texture. The base material color is not used whenever TexFace is enabled.

The top left ball has not had any vertex painting, and the finger is mapped to the middle band, and the texture is mapped to a pink color. As you can see, the base material has VCol Paint and TexFace enabled; the base color blue is not used, but the texture is. With no vertex painting, there is nothing to modulate the UV Texture colors, so the finger shows as white. Transparent areas of the UV Image show as black.

The top right ball has had a pink vertex color applied to the vertical band of faces (in the 3D View window, select the faces in UV Paint mode, switch to Vertex Paint mode, pick a pink color, and Paint->Set Vertex

Colors). The finger is mapped to the middle vertical band of faces, and VCol and TexFace are enabled. The texture is mapped to Alpha black and multiplies the base material alpha value which is 1.0. Thus, white areas of the texture are 1.0, and 1.0 times 1.0 is 1.0 (last time I checked, at least), so that area is opaque and shows. Black areas of the procedural texture, 0.0, multiply the base material to be transparent. As you can see, the unmapped faces (left and right sides of the ball) show the vertex paint (none, which is gray) and the painted ones show pink, and the middle stripe that is both painted and mapped change the white UV Texture areas to pink. Where the procedural texture says to make the object transparent, the green background shows through. Transparent areas of the UV Texture insist on rendering black.

The bottom ball uses multiple materials. Most of the ball (all faces except the middle band) is a base material that does not have TexFace (nor Vertex Color Paint - VCol Paint) enabled. Without it enabled, the base blue material color shows and the pink color texture is mixed on top. The middle band is assigned a new material (2 Mat 2) that *does* have vertex paint and TexFace enabled. The middle band of faces were vertex painted yellow, so the white parts of the finger are yellow. Where the pink texture runs over the UV texture, the mixed color changes to green, since pink and yellow make a green.

If you want the two images to show through one another, and mix together, you need to use Alpha. The base material can have an image texture with an Alpha setting, allowing the underlying UV Texture to show through.

To overlay multiple UV images, you have several options:

- Create multiple UV Textures which map the same, and then use different images (with Alpha) and blender will overlay them automatically.
- Use the Composite Nodes to combine the two images via the AlphaOver node, creating and saving the composite image. Open that composited image as the UV Texture.
- Use an external paint program to alpha overlay the images and save the file, and load it as the face's UV Texture
- Define two objects, one just inside the other. The inner object would have the base image, and the outer image the overlaid image with a material alpha less than one (1.0).
- Use the Material nodes to combine the two images via the AlphaOver or Mix node, thus creating a third noded material that you use as the material for the face. Using this approach, you will not have to UV map; simply assign the material to the face using the Multiple Materials

UV Textures vs. Procedural Textures

A Material Texture, that has a Map Input of UV, and is an image texture that is mapped to Color, is equivalent to a UV Texture. It provides much more flexibility, because it can be sized and offset, and the degree to which it affects the color of your object can be controlled in the Map To panel. In addition, you can have different images for each texture channel; one for color, one for alpha, one for normals, one for specularity, one for reflectivity, *etc.* Procedural textures, like Clouds, are INCREDIBLY simple and useful for adding realism and details to an image.

UV Texture	**Procedural Texture**
Image maps to precise coordinates on the selected faces of the mesh	Pattern is generated dynamically, and is mapped to the entire mesh (or portion covered by that material)
The Image maps once to a range of mesh faces specifically selected	Maps once to all the faces to which that material is assigned; either the whole mesh or a portion
Image is mapped once to faces.	

	Size XYZ in the MapInput allows tiling the texture many times across faces. Number of times depends on size of mesh
Affect the color and the alpha of the object.	Can also affect normals (bumpiness), reflectivity, emit, displacement, and a dozen other aspects of the mesh's appearance; can even warp or stencil subsequent textures.
Can have many for a mesh	Can be layered, up to 10 textures can be applied, layering on one another. Many mix methods for mixing multiple channels together.
Any Image type (still, video, rendered). Preset test grid available	Many different presents: clouds, wood grain, marble, noise, and even magic.
Provides the UV layout for animated textures	Noise is the only animated procedural texture
Takes very limited graphics memory	Uses no or little memory; instead uses CPU compute power

So, in a sense, a single UV texture for a mesh is simpler but more limited than using multiple textures (mapped to UV coordinates), because they do one specific thing very well: adding image details to a range of faces of a mesh. They work together if the procedural texture maps to the UV coordinates specified in your layout. As discussed earlier, you can map multiple UV textures to different images using the UV Coordinate mapping system in the Map Input panel.

UV/Image Editor Menu

There are several menu items, options and features in the UV/Image Editor window that pertain to using and manipulating the images used in UV Textures. In the View->Properties panel, the Anim and Tile options are for the Blender Game Engine only. Animated UV textures are procedural textures that are mapped to UV coordinates.

Window Header Buttons

Pack
: The button that looks like a little package automatically puts a copy of the image file inside your .blend file when you use Image->Open, automatically packing them all together. Use this option to transport only one file (instead of the .blend and all the .png's used), or to isolate your .blend from any changes that may be happening.

Rotation/Scaling Pivot

- **2D Cursor**; , : As analogous to the Object Mode/ Edit Mode 3D Cursor pivot. This will rotate the current UV selection around the 2D cursor position. LMB click to position the 2D cursor at desired location.
- **Median Point**; â § Shift, : As analogous to the Edit Mode Median Point pivot. Enabling the Median Point pivot will calculate a rotation point for a UV selection accordingly to the vertice count or weight of objects. Blender supposes every vertex has the same weight.
- **Bounding Box Center**; . : As analogous to the Edit Mode Bounding Box Center pivot. Blender will encompass your UV selection as tightly as possible with a 2D rectangle or box, and then use it's center as a rotation point.

Sync UV and Mesh Selection
: This will show all UV faces of your Unwrap/UV Map in the UV/Image Window, not only the selected faces. It is very useful when working (applying transformations) on your UV Map layout, or

simply for a better perspective of how faces are linked and laying flat in the UV Map.

UV Vertex select mode
: Enables selection of vertice.

UV Face select mode
: Enables selection of faces.

Sticky UV selection

- **Shared Vertex** ; CtrlC : When welding or stitching, you want to be sure that you sew the UV maps together correctly. At a seam, when unwrapping, two UV coordinates were created, one for each side of the seam. With Stick UVs enabled, a RMB click will not only select the UV vertex closest to the mouse cursor, but also all the other UV vertices that correspond to the same mesh vertex (but who are now shown on the 'other' side of the seam in another UV map).
- **Shared Location** ; AltC : Works in the same way, but only on the UVs that are 'connected', meaning they are within a 5 pixel range of the first selected UV.
- **Disable** ; â § ShiftC : Disables sticky selection.

Snap while CTRL is held during transformation ; â § Shiftâ Tab
: When applying a transformation to a UV selection, hold down Ctrl to snap the selection to your mouse pointer, according to the selected snap method chosen (described below). Once satisfied with position, LMB click to apply transformation or RMB click to cancel.

- **Median** : Move the median of the selection to the snap target (mouse pointer).
- **Center** : Move the current transformation center to the snap target. Can be used in conjunction with 2D Cursor to snap with offset.
- **Closest** : Move the closest point of the selection to the snap target.

Texture Painting
: The magic pencil button changes your mouse and keyboard into a mini-paint program. Paint using the LMB. See Doc:Manual/Textures/UV/Painting_the_Texture

Draw With Alpha
: UV Textures do not have to be totally opaque; they can be *partially transparent, like a partial reflection on a window. Turning on this button makes the display show only the opaque parts of the image.

Draw Alpha Only
: The dot button disregards colors and shows the alpha channel of the image as a BW gradient, with white being Alpha of 1.0 which is totally opaque and black, which is totally transparent. Use this option to see what parts of the object will be transparent based on the UV Texture.

Lock
: When changes happen in this window (UV/Image Editor), other affected windows are updated in real time.

Full Screen
When working on details, remember that you can expand any window to full-screen by pressing â § ShiftSpace to toggle the active window between a pane and full-screen. You can also use Ctrl↑ and Ctrl↓

Image Menu

This menu gives you options when working with the image that is mapped to the mesh.

- **Realtime Texture Mapping**: sets display updates to either UV Coordinates (default) or Reflection. With reflection, the texture shown is like looking in a mirror at the image. Use this when rendering a scene featuring a mirror that has a reflection 'out the window' when there is no 'outside'. This is

sometimes also called using the UV Layout as a *Reflection Map*, namely mapping the image as if it was a reflection. The image you want to use is what would be reflected, namely the view from the reverse camera angle.

- **Texture Painting**: Enables and turns on Doc:Manual/Textures/UV/Painting_the_Texture.
- **Pack Image**: When selected, takes all the images in use and puts a copy inside the .blend file. Use this when sending off your file or packing up for the weekend.
- **Reload**: refreshes the image in Blender by re-reading the source image file. Use this if the artist has updated the image with changes.
- **Replace**: Replaces the Image with a new one that has a different name/location, keeping the UV mapping. The old one is discarded from memory/pack. Use this if you mistakenly opened the wrong file.
- **Save As**, **Open**, and **New**: saves the current image, opens an existing file, and creates a new image, respectively. But I bet you already figured that out.

Thumbnails
Blender has a built-in picture window, the Image Browser window type, that allows you to scroll through directories on your hard disk, and shows you thumbnails only of image files in the directory. When you hover over a file, the header shows you the size and format of the file. Very Handy...even better than the Windoze file browser.

UVs Menu (Image Related)

Sometimes it is necessary to move image files to a new location on your hard disk. Use the Find Image Target Paths script to update the image links. You can fill in the top level root directory name, and Blender will search down from there to find a similar file.

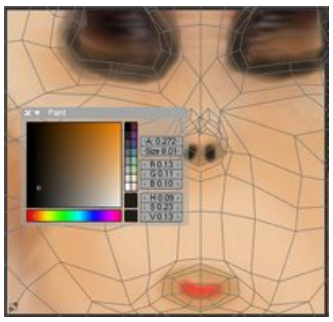
A UV Texture is a picture (image, sequence or movie) that is used to color the surface of a mesh. The UV Texture is mapped to the mesh through one or more UV maps. There are three ways to establish the image used by the UV Texture.

- Paint a flat image in the UV/Image Editor onto the currently select UV Texture, using its UV map to transfer the colors to the faces of the mesh
- Paint the mesh in the 3D View, and let Blender use the currently selected UV map to update the UV Texture ("Projection Painting")
- Use any image-editing (paint) program to create an image. In the UV/Image Editor, select the UV Texture and load the image. Blender will then use that texture's UV map to to transfer the colors to the faces of the mesh

Blender features a built-in paint mode, called Texture Paint, designed specifically to help you edit your UV Textures and Images quickly and easily in either the UV/Image Editor window or the 3D View window. Since a UV Texture is just a special-purpose image, you can also use any outside paint program. For example, Gimp is a full-featured image manipulation program that is also open-source.

Since a mesh can have layers of UV Textures, there may be many images that color the mesh. However, each UV Texture only has one image.

Using Blender's Texture Paint



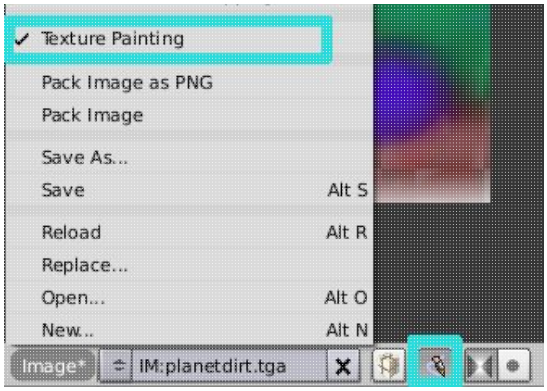
The "Paint" tool in action.
Texture Paint works in both a 3D window and the UV/Image Editor window. In the 3D window in Texture Paint mode, you paint directly on the mesh. In the UV/Image Editor window, you paint on a flat canvas that is wrapped around the mesh using UV coordinates. Any changes made in the UV/Image Editor window show up immediately in the 3D window, and vice-versa.

A full complement of brushes and colors can be selected from a floating Image Paint panel in the UV/Image Editor, or a Paint panel in the Buttons window, Editing (F9) context. Brush changes made in either panel are immediately reflected in the other panel.

When satisfied or at intermittent intervals, please save your image using the UV/Image Editor window. How to use Texture paint is explained in this section.

Square Power of 2
Texture paint is very fast and responsive when working in the 3D window and when your image is sized as a square, power of two; 256x256, 512x512, 1024x1024, etc.

Getting Started



Once you have un-wrapped your model to a UV Texture (as explained in previous pages), you have to:

- Load an image into the UV/Image Editor (Image->Open->select file) , or
- Create a new image (Image->New->specify size) **and** saved it to a file (Image->Save->specify file).

You cannot paint on a mesh in Texture Paint mode without **first** unwrapping your mesh, **and** doing one of the above steps. After you have done these two things, you can modify the image using the Texture Paint mode. Once you do:

- In the 3D View window, select Texture Paint mode from the mode selector in the window header, and you can paint directly onto the mesh.
- In the UV/Image Editor window, enable Texture Painting in the Image menu (shown to the right).
- In the UV/Image Editor header, click the magic pencil (highlighted to the right).

At this time, you may choose to show the alpha (transparency) channel by clicking the button to the right of the magic pencil. The dot icon to the right of that button allows you to paint the alpha channel by itself.

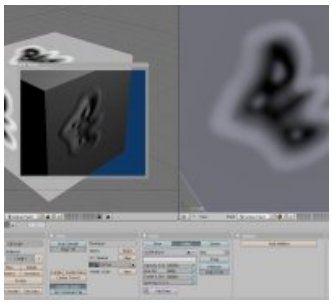
Once you enable Texture Painting, your mouse becomes a brush. To work with the UV layout (for example, to move coordinates) you must disable Texture Painting. To work with the mesh in the 3D View (for example, to move it in 3D space), or select other objects, you must leave Texture Paint mode.

When you enable Texture Painting, use the View->Paint Tool option in the UV/Image Editor window or the Paint panel in the Buttons window to modify paint settings.

All painting you perform in either window will be instantly reflected in the other window (if the 3D View is in *textured* viewport mode). However, the modified texture will not be saved until you explicitly do so by Image->Save in the UV/Image Editor window.

See Outline: If you want to paint directly on the mesh in the 3D View window, change to UV/Face Select Mode first to show the edge outline of the mesh, *then* switch to Texture paint mode, which visually overlays the previous mode, but keeps the outline.
As soon as you enable Texture Painting or switch to Texture Paint mode, a Paint Panel becomes available in the Editing (F9) buttons. This panel has all the same controls as those available in the Paint Tool discussed below. Use this panel if you are only working in 3D view in order to change brushes (colors, patterns, function).

Instant Feedback



If your texture is already used to color, bump map, displace, alpha-transparent, etc. a surface of a model in your scene (in other techie words, is mapped to some aspect of a texture via texture channel using UV as a map input), you can see the effects of your painting in the context of your scene as you paint.

To do this, set up side-by-side windows, one window in 3D View set to Textured display mode, and the second UV/Image Editor window loaded with your image. Position the 3D View to show the object that is UV mapped to the loaded image. Open a Preview window (see 3D View Options for more info) and position it over the object. In the image to the right, the texture being painted is mapped to "Nor" or Normal, and is called "bump mapping", where the gray scale makes the flat surface appear bumpy. See Texture Mapping Output for more information on bump mapping.

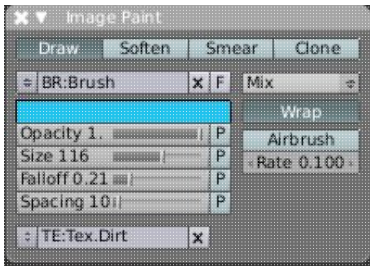
Brushes

To paint on the mesh, you just click and drag the left mouse button across the mesh in the 3D View window, or across the image in the UV/Image Editor window. What that mouse action does to the color of the mesh depends on the Brush you are using. Some brushes add color, some take it away, and some even make your mesh transparent!

Mode: UV/Image Editor

Hotkey: C for Color

Menu: View -> Paint Tool...

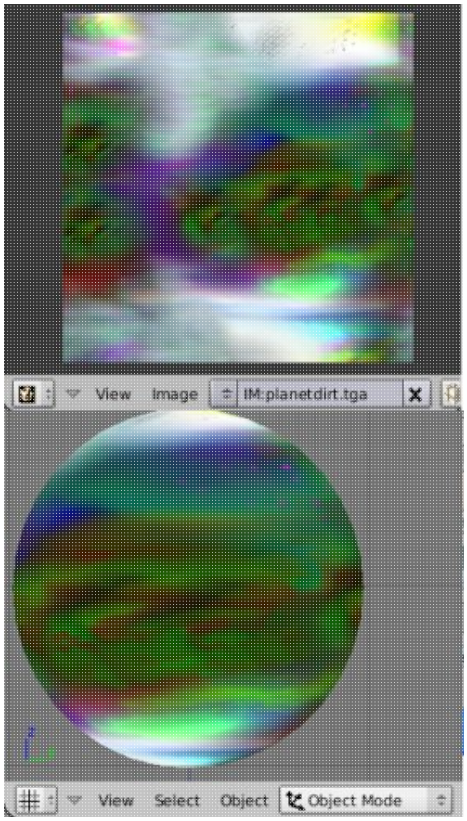


Press C in the UV/Image Editor to pop up the Image Paint panel. With this panel, you can create many brushes, each with unique settings (such as color and width). Use the Brush selector to switch between brushes, or create a new brush. When you add a brush, the new brush is a clone of the current one. You would then change the setting for the new brush. Texture paint has an unlimited number of brushes, and unique user-defined controls for those brushes, set in the Paint Tool panel. You can paint in either window; you can paint right on your mesh in the 3D window, much like vertex painting, or flat on the UV Image in the UV/Image Editor window.

To use a brush, click on its name from the BR: selector up/down arrow. Name your brush by clicking on the name field and entering any name you wish, such as "Red Air" for a red airbrush. To toss out a brush, click the brush delete X button next to its name. If you want to keep this brush around for the next time you run Blender, click the Fake user button next to the brush delete X button.

If you have a tablet pen with pressure sensitivity, toggle the small "P" button next to the opacity, size, falloff and spacing buttons to control these parameters using the pressure of the pen. Using your pen's eraser end will toggle on the Erase Alpha mode.

Brush Mode Options



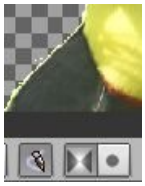
Planet texture example using various brush modes
Across the top and down the side are several controls that determine what the brush does when it draws:

- Draw: The normal brush; paints a swath of color
- Soften: blends edges between two colors
- Smear: when you click, takes the colors under the cursor, and blends them in the direction you move the mouse. Similar to the "smudge" tool of *The Gimp*.
- Clone: copies the colors from the image specified (Tex.Dirt in the example), to the active image. The background image is shown when this brush is selected; use the Blend slider to control how prominent the background image is.
- Wrap: wraps your paint to the other side of the image as your brush moves off the OTHER side of the canvas (any side, top/bottom, left/right). Very handy for making seamless textures.

Mix Mode Options

Mix Mode determines what the brush action does to the image:

- Mix: the brush color is mixed in with existing colors
- Add: the brush color is added to the existing color; green added to red gives yellow.
- Subtract: brush color is subtracted; painting blue on purple gives red
- Multiply: the RGB value of the base is multiplied by the brush color
- Lighten: the RGB value base value color is increased by the brush color
- Darken: tones down the colors
- Erase Alpha: makes the image transparent where painted, allowing background colors and lower-level textures to show through. As you 'paint', the false checkerboard background will be revealed.
- Add Alpha: makes the image more opaque where painted



In order to see the effects of the Erase and Add Alpha mix modes in the UV/Image Editor, you must enable alpha channel display by clicking the Display Alpha or the Alpha-Only button. Transparent (no alpha) areas will show an checker background; the clarity of the checker background depends on the opacity of the image.

Brush Control

The controls for each brush allow you to choose:

- Color: The current brush color shown in the color swatch. Click the swatch to choose a color for the current brush; then pick a new color from the popup color picker applet. You can pick any hue and then a saturation/value, pre-defined colors, or even use the eyedropper to take a sample from anywhere within the Blender window.
- Opacity: how thickly you apply the paint when you click down, like how hard you press down on a spray can button, or how much paint you put on your brush.
- Size: How big your brush or spray area is, in pixels.
- Falloff: How spread-out the spray is. Decrease for a softer brush; increase for a harder one.
- Spacing: The distance between paint spurts when dragging the mouse, relative to the size of the brush.
- Airbrush: When enabled, the brush paints continually while the LMB is held down, rather than only when the mouse is moving.
- Rate: How fast the paint is applied by the Airbrush

RMB click on any part of the image to sample that color and set it as the brush color.

Brush Pattern/Texture



Use the texture selector at the bottom of the paint panel to select a pre-loaded image or procedural texture to use as your brush pattern. Note that in order to use it, you must have a placeholder material defined, and that particular texture defined using the Material and Texture buttons. It is not necessary to have that material or texture applied to any mesh anywhere; it must only be defined. The example to the right shows the effects of painting with a flat (banded) wood texture. Switching the texture to Rings makes a target/flower type of brush painting pattern.

Note: In Clone paint mode, this field changes to indicate the picture image or texture that you are cloning from.

Saving

If the header menu item Image has an asterisk next to it, it means that the image has been changed, but not saved. Use the Image->Save Image option in the to save your work with a different name or overwrite the original image.

UV images
Since images used as UV Textures are functionally different from other images, you should keep them in a directory separate from other images.
The image format for saving is independent of the format for rendering. The format for saving a UV image is selected in the header of the Save Image window, and defaults to Targa (.TGA).

If Packing is enabled in the window header, or if you manually Image->Pack Image, saving your images to a separate file is not necessary.

Options

The Paint controls are also shown on a Paint panel in the buttons window, Editing set.

Mode: Texture Paint

Hotkey: F for Face-select submode

Face Painting: By default, you will be painting the whole mesh and it will look solid. RMB clicking will pick up the color from the mesh at that point. To only paint part of the mesh, press F in 3D View to highlight the UV/Face edges. The faces of the mesh will be outlined for you using the red-green edges in the 3D View, and if you have Shadow Mesh enabled in the UV/Image Editor, the UV Layout will also be shown as a gray outline in the UV/Image Editor window. Now, RMB will select a face, and â § ShiftRMB will select multiple faces. To only paint part of the mesh, press H to Hide the selected Faces when in Texture Paint mode. Just like when editing a mesh, doing a Border select with the LMB button selects only certain faces. Doing a Border select with the RMB button excludes those faces from the selected set of faces.

Selecting some faces and pressing H Hides them from view and thus painting. It also continues to hide them during textured viewport shading in Object view. During your painting session with F active, or in UV Face Select Mode, AltH un-hides them. Unfortunately, H and AltH don't do anything when you are not in this face-select sub-mode (because there aren't faces 'selected').

Using an External Image Editor

If you use an external program to edit your UV Texture, you must:

1. run that paint program (Gimp, Photoshop, etc.)
2. load the image or create a new one
3. change the image, and
4. re-save it within that program.
5. Back in Blender, you reload the image in the UV/Image Editor window.

You want to use an external program if you have teams of people using different programs that are developing the UV textures, or if you want to apply any special effects that Texture Paint does not feature, or if you are much more familiar with your favorite paint program.

WARNING
This is the copy of the Release notes page about Projection Paint released in BLender 2.49, a proper manual page will be written laterâ←¦

2.49 Release

Projection Painting

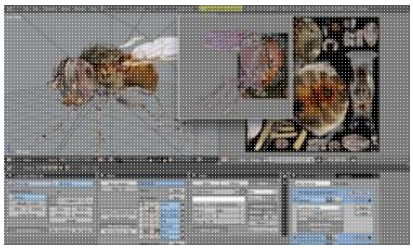


image courtesy of sebastian_k
Painting in the 3D view has been improved to support **projection painting**, which allows you to paint directly onto your model without having to worry nearly so much (or perhaps, at all) about "seams" or "unwrapping manually."

This is useful for anyone making textured models for games (realtime display), painting stencil maps for use with rendering, and for people who just don't want to manually unwrap their meshes.

The essential notion here is as intuitively-simple as it is powerful. From the particular point-of-view represented by a given viewport that you are "looking through," you "paint" the pixels from some currently-selected image "onto" the model. Blender automatically determines the corresponding UV-coordinate locations for you, selecting the proper pixels from the source UV-map and using them to modify the pixels in the target UV-map while showing you the effects of your work in real time. You can rotate the object at any time, select different source images at any time, and adjust many other characteristics of how pixels are selected, mixed and applied. It really *is* "like painting on the model itself."

Projection painting works on a UV-mapped texture. Thus, you need a non-overlapping set of UV coordinates as well as an image to paint to. To automatically generate an UV set suitable for painting, you can just unwrap your mesh with the "Lightmap UV pack" option or "Smart Projections."

Note: Of course you *can* use any UV mapping that you like, prepared any way you like. You should just make sure that the map doesn't have any overlapping areas, lest your "paint" wind up in several places on the model's surface at once. Projection painting is sometimes used to "touch up" image-maps that were prepared using other methods; or, vice-versa.

Also See...

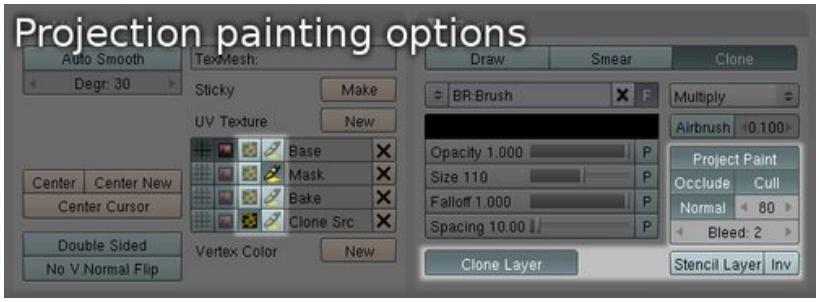
The player will show in this paragraph

Features



image courtesy of ropsta

- Screen resolution independent
 Detail will remain the same no matter how far you're zoomed in
- Blending modes, Texture Brushes (Nodal brush textures)
- Tablet pressure sensitivity
- Face select
 Only paint onto selected faces (FKey toggles)
- Paints directly into textures
 Allowing you to paint into diffuse, specular maps and other textures used by GLSL shaders
- Paint onto the modified mesh
 Subsurf, Armatures, Multires etc. can be enabled while painting

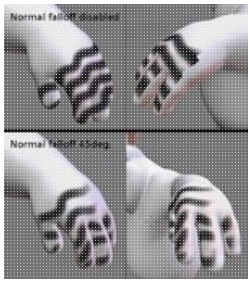


New buttons to select clone and stencil UV layers as well as 3D paint options (right)



Occlude disabled
Geometry Occlusion

When this is enabled only exposed pixels are affected, this also allows for 3D stencils to be used to mask out areas of the surface too. This is normally enabled along with the *Cull* and *Normal* settings.

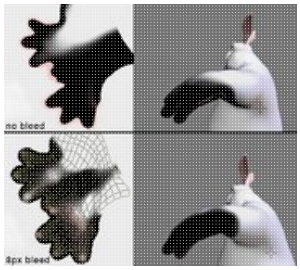


normal falloff on/off
View Normal Falloff

As faces point away from the view the brush strokes fade away to prevent harsh edges or brush smearing (optional).

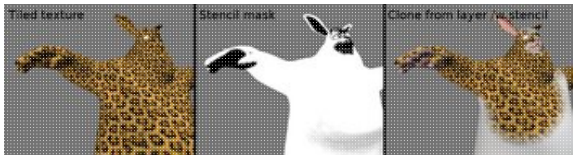
The normal angle can be adjusted to only allow painting onto very front facing areas of the mesh.

When painting onto models with large faces and sharp angles (low poly building for example) you may want to disable this option.



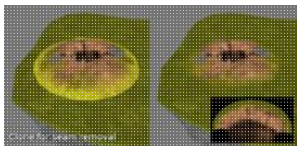
8px bleed comparison
Seam Bleed

Extend the paint beyond UV island bounds to avoid visual artifacts (like bleed for baking)



Clone between layers with a stencil
Clone between UV Layers

Paint from 1 UV layer to the current one, useful for taking photos from multiple angles and painting them onto an unwrapped texture



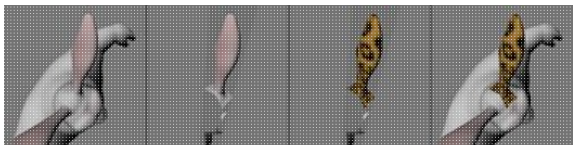
Clone with an offset to remove seams
Clone with offset

Useful for painting over UV seams which isn't possible to do in a 2D image editor, otherwise its similar to cloning in the Gimp. (Ctrl+LMB to set the 3D cursor location to clone from)

File:ProjectPaint stencil simple.png

Stencil (mask)
UV layer stencil mask

Use another UV layer as a mask for painting, (like quickmask in The Gimp)



☐
Clip the view by a bounding box
View clipping

Painting can be constrained to a bounding box (Alt+B to set view clipping).

This helps to avoid painting onto overlapping geometry.

File:ProjectPaint opacity.png

Limit opacity
Correct Opacity Blending

never accumulate over the opacity set for a brush (except when Airbrush is enabled)

Implementation

- screen to UV Image mapping initialization as needed.
- multi-threaded painting and mapping initialization.
- floating point as well as 8bit channels.
- partial OpenGL updates *(avoid sending the entire image to the graphics card)*
- tile based undo *(only store modified tiles for each undo step)*

Limitations

- Using textures with dimensions other than 128, 256, 512, 1024, 2048, 4096 will be a lot slower to refresh.
- Large faces with a single vertex behind the view in perspective mode will be ignored.
- Overlapping UVs will show artifacts when painting onto the same pixels from different faces (only when air-brush is disabled).
 This includes Mirrored UVs in some cases.
- Hard edges with the edge split modifier are detected as UV island boundaries when bleed is enabled.

Future Plans

- Easy image stamp option
- Normal map painting options
- Texture mapping modes (world, orco, normal, reflection)
- Project from face normal (even width line over a curved surface)
- Auto select paint depth exclude distant geometry.
- Mirror painting
- Keep screen:uv cache between brush strokes, generate while idle. (will improve user experience with high-res textures)

Brush System

- thumbnails to select brushes *(like in the gimp)*
- Curve falloff (like sculpt currently has)
- Options jitter, paint image sequence.

World options

Blender provides a number of very interesting settings to complete your renderings by adding a nice background, and some interesting 'depth' effects. These are accessible via the Shading Context (F5) and World Buttons sub-context () shown in *World Buttons*. By default a very plain uniform world is present. You can edit it or add a new World.



World Buttons
You have:

- Background
- Mist
- Stars

Note
Some of the settings under the World panel in Blender affect lighting so you find them under the Lighting chapter (see Ambient Light, Exposure and Ambient Occlusion), while some Physics settings are related to the Game Engine, so be careful don't get confused!

World Background

Mode: All Modes

Panel: Shading/World Context → Preview

Hotkey: F8

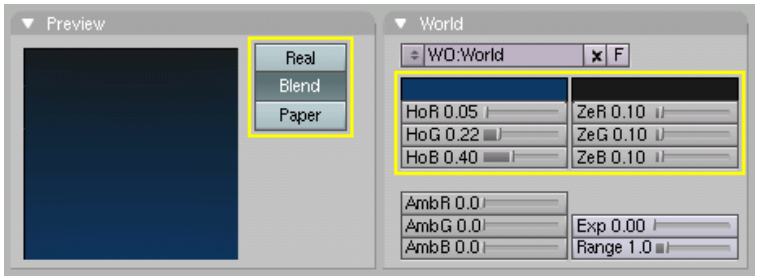
Description

The world buttons let you set up the shading of your scene in general. It can provide ambient colour, and special effects such as mist, but a very common use of a World is to shade a background colour.

Background Image in Render
To use an image as your render background, see BackBuf images specified in the Output Panel

Background Image in 3D
To use an image as a background image in your 3D view, for example as a reference when doing a model, see using a Background Image

Options



Background colors

HoR, HoG, HoB
: The RGB color at the horizon

ZeR, ZeG, ZeB
: The RGB color at the zenith (overhead)

These colors are interpreted differently, according to the Buttons in the Preview Panel (*Background colors*):

None Enabled
: If none of these three buttons is checked, your background will just be plain flat color (using the horizon one).

Blend
: The background color is blended from horizon to zenith. If only this button is pressed, the gradient runs from the bottom to the top of the rendered image regardless of the camera orientation.

Real
: If this option is added, the gradient produced has two transitions, from nadir (same color as zenith) to horizon to zenith; the blending is also dependent on the camera orientation, which makes it more realistic. The horizon color is exactly at the horizon (on the x-y plane), and the zenith color is used for

points vertically above and below the camera.

Paper
: If this option is added, the gradient keeps its characteristics, but it is clipped in the image (it stays on a horizontal plane (parallel to x-y plane): what ever the angle of the camera may be, the horizon is always at the middle of the image).

Textures

Instead of a color, or blend of two colors, Blender can use an 2D image which it maps to a very large Box or sphere which encompasses the entire scene, or which it maps to a virtual space around the scene.



Texture buttons
The World Buttons also provide a stack of texture channels. They are used much like the Materials textures, except for a couple of differences (*Texture buttons*). The textures can be mapped according to:

View
: The default orientation, aligned with the co-ordinates of the final render

AngMap
: Used to wrap a standard hemisphere angular map around the scene in a dome. This can be used for image based lighting with Ambient Occlusion set to sky color. You'll generally need a high dynamic range image (HDRI) angular map (the look like a weird spherical image).

Sphere
: Sphere mapping, similar to that of materials

Tube
: Wrap the rectangular texture around in a cylinder, similar to that of materials

Object
: Position the texture relative to a specified object's local texture space

The texture affects color only, but in four different ways:

Blend
: Makes the Horizon color appear where the texture is non-zero

Hori
: Affect the color of the horizon

ZenUp
: Affect the zenith color overhead

ZenDo
: Affect the zenith color underneath

Mist

Mode: All Modes

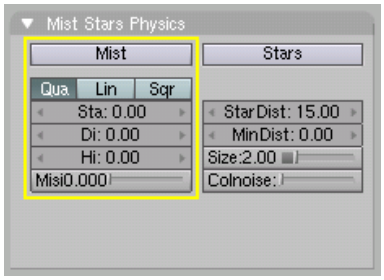
Panel: Shading/World Context → Mist Stars Physics

Hotkey: F8

Description

Mist can greatly enhance the illusion of depth in your rendering. To create mist, Blender makes objects farther away more transparent (decreasing their Alpha value) so that they mix more of the the background color with the object color. With Mist enabled, the further the object is away from the camera the less it's alpha value will be.

Option



Mist buttons

Mist
: Toggles mist on and off.

Qua / Lin / Sqr
: The decay rate of the mist, Quadratic, Linear, and Square Root. These settings the rate of change of the mist's strength further and further into the distance.

Sta
: The distance from the camera at which the mist starts to fade in

Di

The distance from the Start of the mist, that it fades in over. Objects further from the camera than Sta+Di are completely hidden by the mist.

Mist distances
To visualize the mist distances in the 3D View, select your camera, go to Editing Context (F9) and enable *Show Mist* in the *Camera* Panel.

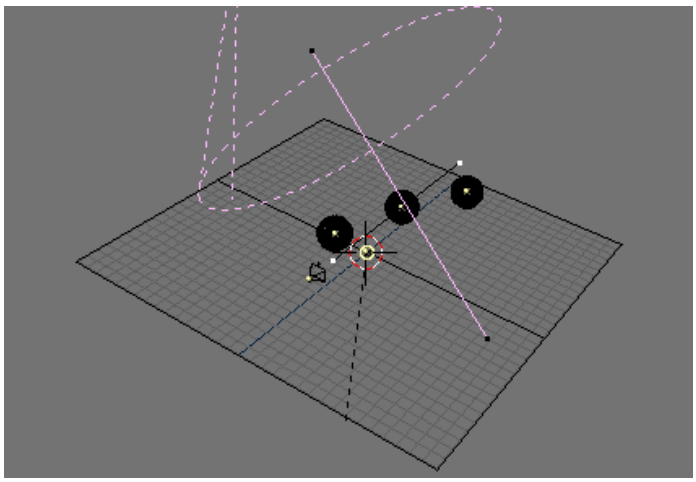
The camera will show mist limits as a line projecting from the camera starting from *Sta* and of distance *Di*.

Hi
: Makes the mist intensity decrease with height, for a more realistic effect. If greater than 0, it sets, in Blender units, an interval around z=0 in which the mist goes from maximum intensity (below) to zero (above).

Misi
: An overall intensity, or strength, of the mist.

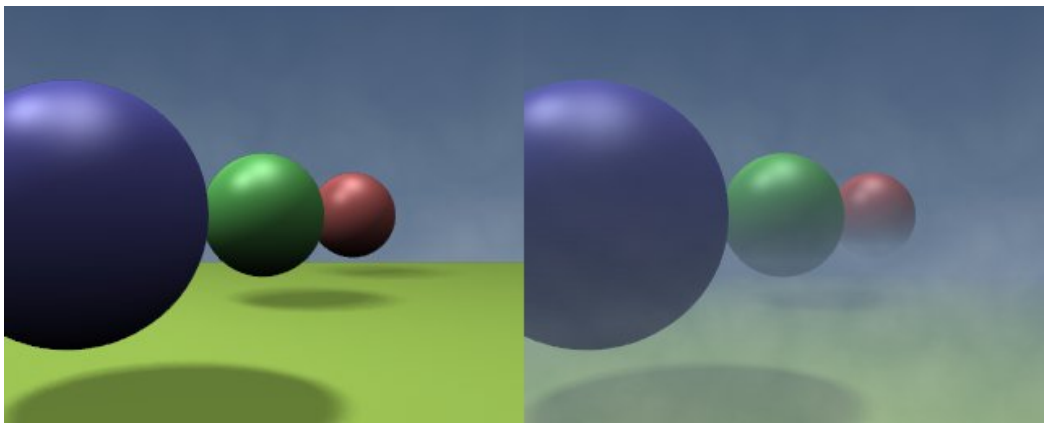
Examples

Mist test setup shows a possible test set up.



Mist test setup

Rendering without mist (left) and with mist (right). shows the results with and without mist. The texture is a plain procedural cloud texture with Hard noise set.



Rendering without mist (left) and with mist (right).

Stars

Mode: All Modes

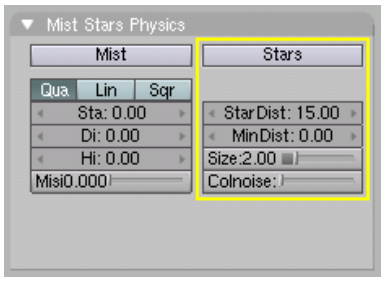
Panel: Shading/World Context → Mist Stars Physics

Hotkey: F8

Description

Stars are randomly placed halo-like objects which appear in the background. The Stars settings are on the right-hand side of the Mist Stars Physics Panel (*Star buttons*).

Options



Star buttons

StarDist
: The *average* distance between stars. Stars are intrinsically a 3D feature, they are placed in space, not on the image!

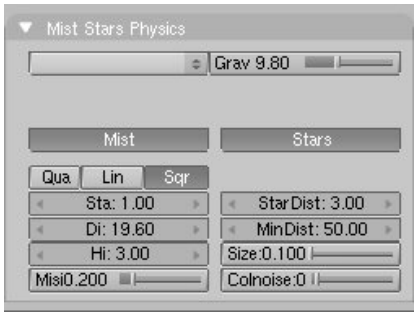
MinDist
: The *minimum* distance from the camera at which stars are placed. This should be greater than the distance from the camera to the *furthest* object in your scene, unless you want to risk having stars *in front* of your objects.

Size
: The actual size of the star halo. It is better to keep it much smaller than the proposed default, to keep the material smaller than pixel-size and have pin-point stars. This is much more realistic.

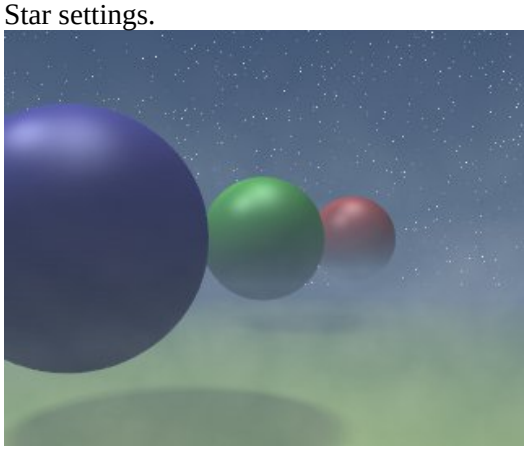
Colnoise
: Adds a random hue to the otherwise plain white stars. It is usually a good idea to add a little ColNoise.

Examples

Star rendering. shows the same misty image from the Mist chapter but with stars added. The Stars settings used for the image are shown in *Star settings.*.



Star settings.



Star rendering.

Animation is making an object move or change shape over time. Objects can be animated in many ways. They can be animated as one or more simultaneous ways:

Moving as a whole object
: Changing their position, orientation or size in time;

Deforming them
: animating their vertices or control points;

Character Animation via Armature
: animated to deform by the movement of bones inside the mesh, a very complex and flexible interaction that makes character-shaped objects appear to walk and jump.

In this chapter we will cover the first case, but the basics given here are actually vital for understanding the following chapters as well.

Three methods are normally used in animation software to make a 3D object move:

Key frames
: Complete positions are saved for units of time (frames). An animation is created by interpolating an object fluidly through the frames. The advantage of this method is that it allows you to work with clearly visualized units. The animator can work from one position to the next and can change previously created positions, or move them in time.

Motion Curves
: Curves can be drawn for each XYZ component for location, rotation, and size. These form the graphs for the movement, with time set out horizontally and the value set out vertically. The advantage of this method is that it gives you precise control over the results of the movement.

Path
: A curve is drawn in 3D space, and the Object is constrained to follow it according to a given time function of the position along the path.

The first two systems in Blender are completely integrated in a single one, the *IPO* (InterPOlation) system. Fundamentally, the IPO system consists of standard motion curves. A simple press of a button changes the

IPO to a key system, without conversion, and with no change to the results. The user can work any way he chooses to with the keys, switching to motion curves and back again, in whatever way produces the best result or satisfies the user's preferences. The IPO system also has relevant implication in Path animations.

Chapters

General Principles and Tools
- Ipo Types
- Creating Ipo Keyframes
- Editing Ipo Curves and Keyframes
- The Timeline - The Timeline Window

Animation by Moving Objects Around
- Moving objects on a Path
- Changing Object Layers

Animation by Basic Deformation
- Deformation
 - Shape Keys
 - Absolute Shape Keys
 - Deforming by a Lattice

See also Hooks - Uses a modifier as a way to change the shape of a mesh. Sorta like sticking a fish hook in a mesh and pulling. Uses the principles discussed in Shape Keys.

The BSoD Introduction to Character Animation tutorial is a good starting point for learning character animation. Even if you never used Blender before.

Interpolation (Ipo) is the process of estimating an object's position (or other attributes) based on a known start and end value, the time between the start and the end, and a continuous curve that connects the two points. If I stand in 3D space (0,0,0) at frame 1, and I want to move to (2,2,0) by frame 10, and I define a curve of acceleration, then Blender can interpolate, for frame 5, that I will be half-way there, or at (1,1,0). To make that interpolation, you define a bezier curve which connects the two points, and Blender uses the curve to figure out the intermediate values.

Interpolation in Blender

The Interpolation (IPO) animation concept is found and applied to many different things in Blender. It makes no difference whether an object's location is controlled by an IPO, or its material settings, or its size, etc. When an IPO is associated with an object in some type of relationship, the IPO controls one or more aspects of the object at any particular frame in the animation. For example, there is an Object type of IPO, that, when assocated with a specific Object, can control that object's location in the XYZ space, its rotation, its scale, and so on. Exactly which aspect is controlled depends on which *channels* are identified. For example, there is a Location channel that you can select when you key an object in 3D View. That Location channel is represented by three IPO curves: LocX, LocY, and LocZ that specify the object's (X, Y, Z) location. Once you have learned to work with object IPOs, how you work with other channels will become obvious. Blender does distinguish between different *types* of IPOs and the channels that are available, and keeps track of all of that automatically.

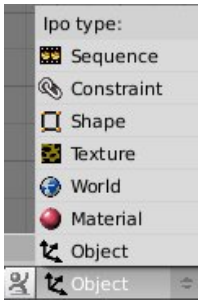
Every type of IPO block has a fixed number of available *channels*. These each have a name (LocX, SizeZ, etc.) that indicates how they are applied. When you add an IPO Curve to a channel, animation begins immediately. At your discretion (and there are separate channels for this), a curve can be linked directly to a value (LocX...), or it can, in some instances, affect a variance of it (dLocX...). The latter enables you to move an object as you would usually do, with the Grabber, without disrupting the IPO. The actual location is then

determined by IPO Curves *relative* to that location.

The Blender interface offers many options for copying IPOs, linking IPOs to more than one object (one IPO can animate multiple objects.), or deleting IPO links. The IPO Window Reference section gives a detailed description of how to work with curves and their handles.

If there is an IPO curve for a channel, the interpolated value overrides whatever panel setting is changed by you. You can change the property value through the panel and re-key (through the I-key command) to take on that value for that frame. Otherwise, it will reset to the interpolated curve value when you render or change frames! This can be very confusing if you forget that a channel has been automated. You can always select a channel, select the curve, and Delete X it if you do not need it to be animated.

Ipo Curve Types



Ipo Types menu
As well as the default Object Ipo, there are Ipos to control animation of Materials, Worlds, Shapes, Constraints, Video Sequence strips and more. For every kind of object in Blender, you can animate (change over time):

- Object - the object itself as a whole, such as its location, rotation, layer...
- World - World environment settings, such as horizon color, ambient light...
- Constraint - limits placed on the object
- Sequence - a video effect strip in the Video Sequence Editor.



Ipo Types menu
Depending on the kind of object that is selected, additional animatable aspects become available. For example, the Curve Ipo type appears if the selected object is a curve and not a mesh; the Lamp Ipo type only appears if the selected object is a lamp. Whole other additional channels, or sets of channels become available for animating based on the object selected:

- Mesh - Materials (colors, specularity, alpha), Textures and the Shape of the mesh
- Curve - Shape as well as Path Speed (to vary the speed of an object following the curve as a path)
- Surface - Shape of NURBS surfaces (donuts, spheres) deformed
- Meta - no additional animation
- Text - no additional animation
- Empty - nothing additional to animate
- Camera - Lens, clipping boundaries, Aperture, Focal Distance...

- Lamp - energy (brightness), color, texture...
- Armature - Poses, which are rotation sets for all the bones in the armature
- Lattice - Shape of the Lattice

On the IPO window header, after the menu, there is an eye icon. Enable this to **mute** (or blind Blender to) the channel and any IPO curves. This is handy when debugging your scene in trying to determine, for example, whether lighting or material changes are making your character look funny.

Next over is the IPO Type, and you can click the up/down selector to change between the various IPO Types. For those channels that map to a texture, there will be a texture channel number selector that indicates which texture channel mapping is being automated.

Next over is the name of the IPO datablock. Different materials can share the same datablock. For example, you can make all materials fade out by creating an IPO curve of just the Alpha channel. If you applied that curve to each material in your scene, then regardless of the material's color, spec, etc, they all would fade out as their alpha channel was controlled by the curve. Click the X to delete, or the F to Fake user and retain this curve, whether or not anyone actually uses it.

Next over are the curve copy/paste arrow buttons, the zoom selector (click it and then box select a range of frames to focus on), and the lock icon which forces other windows to update when you make changes.

Animating Materials

Material Ipos can be used for animating a Material so that it changes over time; for example, a character's face can turn from red to blue in 10 seconds as he holds his breath. Just as with objects, IPO Material curves can be used to specify 'key positions' for Materials. Each aspect of a material is animated through a channel. Just as a Object IPO has channels for automating Location X, Location Y, etc, the Material IPO has channels for Color Red, Color Green, etc.

Any animation you do for a material will be reflected in all objects (and their material index) that share that material. Any animation can be driven by some aspect of another object (see IPO Drivers). Once a material channel is animated, the panel sliders and values reflect what the value is for that frame. Changing the value in the panel will not have any effect, because the value is being determined by the interpolated curve.

To access these channels in the IPO window, select Material as the Ipo type. There may be more channels than can be displayed in the window. To see channels that have scrolled off the screen, drag your MMB  to pan the window up and down. Pan the workspace to see the curve (or work in Keyframe mode.)

Creating Channels

There are dozens of Material channels. For memory and processing efficiency, and to just help your workflow efficiency, when you first create an IPO for a material, you get to choose a set of commonly used materials channels. With the mouse in the Buttons Window, displaying the Shading F5 Materials buttons, the command I calls up a pop-up menu with options for various sets of Material channels:

RGB
: Color of the material

Alpha
: Transparency of the mesh/material

Halo Size
: Diameter of the Halo, if selected

Mode

An integer that specifies different settings of the material
- ◊ 1 - Traceable for Ray-tracing, but does not receive shadows
- ◊ 2 - Not traceable, but receives shadows
- ◊ 3 - Shadeless
- ◊ and so on. The easiest way to determine the rest is to play.

All Color
: Many color-related channels - base, specular, alpha, emit (16 in all)

All Mirror
: Many ray-traced reflectivity variables - values and color

Ofs
: The offset of the selected texture to the surface

Size
: The dimensions of the texture applied to the mesh/material

All Mapping
: all the texture settings - color (actual color and amount of influence, Normal amount, offsets, size and so on.

The above-described popup list is just a commonly used set of channels, and not all possible channels are included in the superset of those lists. e.g., the RGB option just initializes/sets a curve for the three R, G, and B channels. For Emit, the "All Color" selection keys that channel among others.

The complete exhaustive list of Material channels is listed down the side of the IPO Window. Choosing a selection from the popup menu may, but not necessarily, set curves and values for every related IPO channel. For example, you can go through all the popup menu selections, and yet Mirror color is not keyed, even though you would think it should be when you select "All Mirror".

Consider memory and processing efficiency when selecting a set. It's inefficient to create a curve (and then have to interpolate it for every frame) when chances are it is not employed/utilized to actually do anything. The Material sets are there to keep the number of channels that Blender has to shift through to a minimum. Remember that when rendering, Blender has to go through every object x materials x channels to determine a value, so more channels is a multiplier effect on render time.

Animating Channel Values

After you select a set, IPO curves are created for the channels in that set, and the current values for those channels are locked in. Not all values for all channels are locked in. The channel may be listed, but not have a curve initialized. There are two ways to initialize a curve for a selected channel. First, LMB select the channel (for example, Emit). Then, in the IPO window, either:

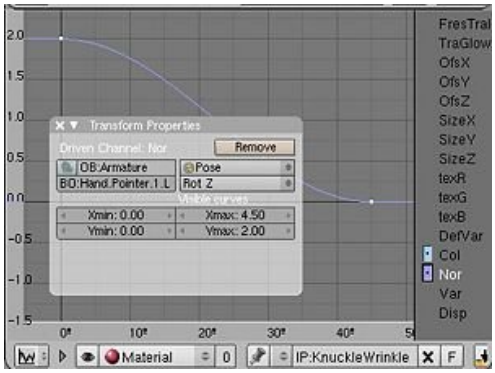
- CtrlLMB click anywhere in the window itself to insert a control point where the cursor is. So if you want emit to be 0 initially, click at (frame 1, value 0). Blender creates a constant value (horizontally flat) curve. Move the mouse/scroll the window to the frame where you want it to change, and click again. Blender creates a curve, which you can then change.
- select another channel and copy that curve to the buffer in the IPO Editor (use the little down arrow on the header). Then select the target channel and apply it to the desired selected channel (up arrow button).

Material curves are like any other curve type, and can be different kinds of control points, curve types, extend modes, and can be displayed as Curves or Keyframes.

Animating Texture Channel Mapping

If you are in a Material, Lamp or World Ipo Block, then a small Number field appears next to the IPO type Menu in the IPO Window toolbar. This indicates which texture channel is active. Recall that you have Map Input and Map To parameters for each texture in the Texture tab of the Materials Buttons Panel. These settings can be animated. The mapping for all texture channels can be controlled with Ipo Curves. Click the number and enter a number to change the texture channel being animated. Channel 0 is the top channel or base texture, and goes from there (channel 1 is the next channel down in the list). The Texture itself can be animated, but that is a different IPO Curve Type.

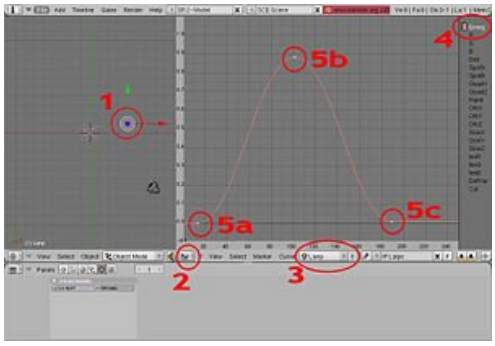
Driving Channel Values



In addition to animating a channel, such as a Material's Color, manually for specific frames, you can drive the change by having some other object do something, like rotate on its Z axis. As that other driver object rotates, blender changes the material color. This relationship is called a Driver relationship, and it works through an Ipo curve as well. Another example is a dimmer knob that rotates to control the lamp's energy channel. If, in the animation, you animate the knob turning, the energy output of the lamp will go up and down. This one knob can control many lamps; for example rotating it in the Z direction could make the inside lights go brighter and the outside sun get darker. the relationship between the knob turning and the energy changing is controlled through a shaped curve, so you can make the relationship smooth, linear, or abrupt.

In the example to the right, we have a finger mesh with a ridged texture on the top texture channel (channel 0). When the bone (Hand.Pointer.1.L - the first bone of the left hand's pointer finger) is straight (rotation Z is 0), the Normal value of the texture channel is 2.0. As the finger bone bends toward 45 degrees, the Nor value decreases to 0, simulating the skin smoothing out. If the bone continues past 45 degrees, the Nor value remains at 0.

Animating Lamps



A Lamp IPO is yet another type of Ipo that controls aspects/channels of a lamp. You can, for example, change the color of a lamp over time by setting Ipo curves for the red, green and blue channels. The example to the right shows how to animate the energy (brightness) of a lamp:

1: Select the Lamp you want to control via RMB in the 3D View.
2: Make sure you have an IPO window open.
3: Set the IPO type to Lamp.
4 Click on "Energy" to establish a lamp energy IPO.
5a,b,and c: Use CtrlLMB click to set the points for your IPO curve.

Animating World Settings



Pressing I in the Buttons window while the material World settings are displayed allows you to make an Ipo curve for any of the channel groups shown to the right. Within each channel group are individual channels.

For example, in the Horizon group there are three channels, HorR, HorG, and HorB for Horizon Red, Horizon Green, and Horizon Blue, respectively. You see these channels by selecting the World Ipo Type in the Ipo Window. Setting up an Ipo curve (via setting the color and pressing I in the Buttons window, or by manually creating points via CtrlLMB clicking in the Ipo window) animates these settings over time. By animating the world settings, you can simulate sunrise to sunset.

Animating Constraints

At the bottom of each Constraint panel is an Influence slider. Next to the slider is a Show and Key button. The Key button creates an Ipo Datablock and keys the influence setting at that frame. You can view the Ipo curve by selecting the Constraint type of curve set. This curve should vary between 0 and 1. A practical example is when you want a camera to track an object for a while, and then perhaps shift attention to a different object. In this example, you would create one constraint, with an influence of 1.0 between frames 10 and 50, and zero at all other times. If you use a bezier curve and key 0 at frame 1, then the camera will swing from its default position to point at the object between frames 1 and 10, track to the object until frame 50, and then resume its normal transformation.

Sharing Ipo Datablocks

Animation curves, such as delta location, can be shared by multiple objects. Defining a relative motion (dLoc) curve for one object would result in an Ipo Datablock that holds that displacement curve. You can easily assign that same curve to another object simply by selecting the object and using the Ipo selector (in the Ipo Window) to choose that Ipo curve. That Ipo datablock will then have 2 users. Each object will, during the same time frame, move relative to their starting position by the same amount.

If you want to delink the selected object from its Ipo curve, click the X next to the Ipo Name. If you want to make some customizations to the motion of the active object only (and not all the other objects that share this Ipo), you have to make a single user copy of the Ipo curve for this object. To do this, click on the Number of Users button (the 2 in the example above) located on the header next to the name. Blender will prompt you to

confirm making a single user copy, and if you confirm, a new Ipo will be created (the name will change) and it will be assigned only to the active object.

See Also

- Reference Manual: IPO Window

Mode: All Modes

Hotkey: I

Menu: Object -> Insert Keyframe



The simplest method for creating an object IPO is with the "Insert key" (I) command in the 3DWindow, with an Object selected. A Pop-up menu provides a wide selection of options (*Insert Key Menu.*). Selecting an option (eg. Loc) saves the current X/Y/Z location and everything takes place automatically:

- If there is no IPO block, a new one is created and linked to the object.
- If there are no IPOCurves in the channels LocX, LocY and LocZ, these are created.
- Vertices are then added in the IPOCurves with the exact values of the object location.

Example

Position your Scene to Frame 1. Add a mesh object to your Blender Scene. The object is added at the 3D Cursor loctation in edit mode. Press â Tab to leave edit mode, and I to Insert a key. In the popup Insert Key window, select Loc. In your Ipo Window, you now see, for the Object mode, the LocX, LocY, and LocZ channels have been filled and selected. There is now a (flat) curve with curve points at Frame 1. The Y value of these points define the location (value) of that channel at that point in time.

Go 30 frames further on (3 x â) and move the object by Grabbing it and dropping it somewhere else. Again use I. Now we can immediately press â µ Enter since Blender remembers our last choice (Loc) and will highlight it. The new position is inserted in the IPO Curves at frame 30.

We can see the animation by slowly paging back through the frames (â). The object moves between the two positions. In this way, you can create the animation by paging through the frames, position by position. Note that the location of the object is *directly* linked to the curves. When you change frames, the IPOs are always re-evaluated and re-applied. You can freely move the object within the same frame, but as soon as you change frame, the object 'jumps' to the position determined by the IPO.

The rotation and size of the object are completely free in this example. They can be changed or animated with the Insert key procedure by selecting from the I menu the other options such as Rotation, Size and any combination of these.

Recording IPOs

Blender has a built-in physics simulation engine ("Bullet") which allows you to define objects as having a mass, and being solid, and having velocity and angular momentum. Once you set this up, you can use the main menu Game->Record Game Physics to IPOs" selection to toggle on recording. Now, when you next press P for Play, the actual location and rotation of the selected object will be recorded for each frame of the animation.

Use this feature for recording complex interactions of multiple animated objects. For more information on using the Blender Game Engine, check out the appropriate wiki book from the main wiki page.

Linking IPOs

An Ipo is (in Blender) a datablock, in that it is a separate object. You make it act on an object by linking it to the object. You can offset that link using the TimeOffset in the Anim settings for the object.

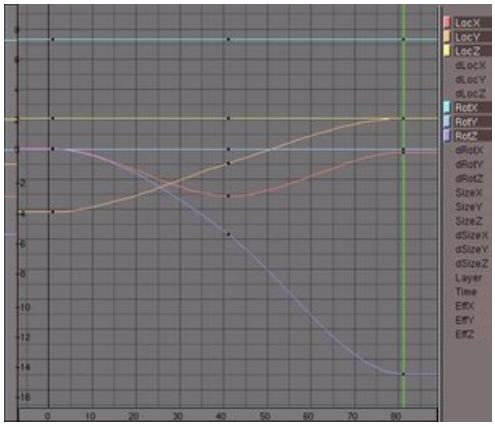
There are two views of your animation, and thus two ways to define and adjust animation: curves and keys. A **curve** connects two points. The shape of the curve is adjusted by moving handles for the points. The handles assert the degree of influence that the point has on the curve. A **key** is a combination of values of points at a particular point in time. Using keys in animation is called Keyframing.

Ipo Curves

Mode: Ipo Curve Editor

Hotkey: C

Menu: View->Show Keys (to disable)



The Ipo Window - Curve display mode
The Ipo Curve Editor allows you to edit the 2D curves that define animation in Blender. The curves represent the edited value in the vertical (Y) axis (location, size, rotation, energy, etc.) and time on the horizontal (X) axis. The rate of change of these values over time can be seen in the slope of the curve.

You can turn any window into an Ipo Curve Editor with the Window Type menu, but it is often handy to have both a 3D View and an Ipo editor at the same time. This shows all the Ipo Curves, the channels used and those available. You can zoom in and out the Ipo Window and translate it just as every other Blender Window.

In addition to the standard channels, which can be set via I, you have the *delta* options, such as dLocX. These channels allow you to assign a *relative* change. This option is primarily used to control multiple objects with the same Ipo. In addition, it is possible to work in animation 'layers'. To set an Ipo Curve for these channels, select the channel and then create control points via CtrlLMB clicking in the Ipo Window.

The Ipo curves have an important feature that distinguishes them from normal curves: it is impossible to place more than one curve segment horizontally. Loops and circles in an Ipo are senseless and ambiguous. An Ipo can only have 1 value at a time. This is automatically detected and corrected in the Ipo Editor. Every Ipo channel that has already been keyed is shown in the channel list to the right of the window with a color button by its side.

Curve showing/hiding

You can show an Ipo curve by using LMB on the relative Ipo channel name on the right (one that already possess a color button next to list channel name). By doing this you hide all previously shown Ipo curves. If you want to show many Ipo curves together (for example if you're working with a group of channels such as the LocX, LocY, LocZ) use â § Shift LMB on every channel name you want to be shown (their names appear in white).

Curve Selection

Mode: All Modes

Menu: Ipo Curve Editor → Select

Some of the standard Blender selection methods work in the Ipo Curve editor. RMB click on a curve to select it. You can also use LMB on the color button next to its name. You can select many curves by using â § ShiftRMB on every curve that is shown or by using â § ShiftLMB on every color button. A curve that is selected shows its control points in white and its respective color button appear outlined in black (if it is the last selected when selecting multi curves the other button appear as pushed).

Options

RMB
: Select an Ipo curve under the mouse pointer

Select/Deselect All A
: Select all visible Ipo curves

Border Select B
: Select Ipo curves within an interactively mouse-drawn box

Curve Manipulation

Mode: All Modes

Menu: Ipo Curve Editor → Curve → Transform

Ipo curves can also be manipulated with many of the standard 2D transformations. The same axis locking options apply here too. Moving curves left and right will move them backwards and forwards in time.

Options

Grab/Move G
: Move the selected Ipo curve

Scale S
: Scale the selected Ipo curve

Curve Interpolation

Mode: All Modes

Hotkey: T

Menu: Ipo Curve Editor → Curve → Interpolation Mode

The interpolation mode determines how the Ipo values are calculated in between each curve point.

Options

Constant
: After each point of the curve, the value remains constant (horizontal). No interpolation takes place.

Linear
: Each curve point is connected by a straight line

Bezier
: The standard smooth curve interpolation

Curve Extend Mode

Mode: All Modes

Hotkey: E

Menu: Ipo Curve Editor → Curve → Extend Mode

The curve extend mode determines how the Ipo values are calculated outside the horizontal limits of the Ipo curve. This can be used to repeat a small section of animation or make animation continue endlessly.

Options

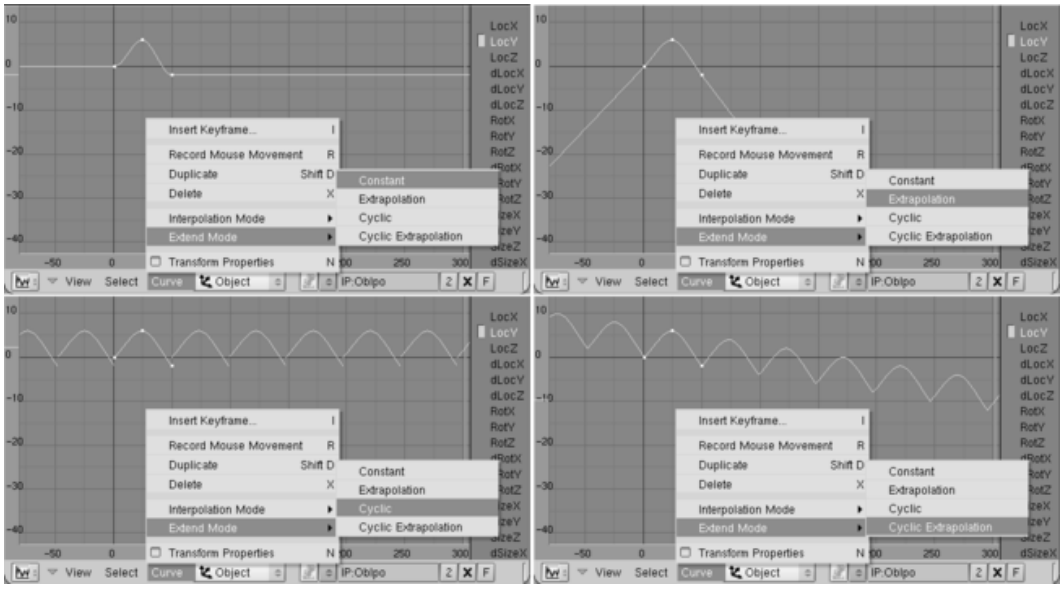
Constant
: The ends of selected curves are continuously (horizontally) extrapolated.

Extrapolation
: The ends of the selected curves continue in the direction in which they ended

Cyclic
: The complete width of the curve is repeated cyclically

Cyclic Extrapolation
: The complete width of the curve is extrapolated cyclically

Examples



Extended IPOs

Editing Curve Points

Mode: All Modes

Hotkey: â Tab

Menu: Ipo Curve Editor → Curve → Edit Selected

As well as manipulating the curves themselves as a whole, you can also edit the individual points that define the curve. Curve points are added when a key frame is inserted (usually with the Insert Key I menu), and these points can be manipulated in time (X axis) and value (Y axis). Hit â Tab to go in edit mode and be able to select each curve point independently. If you select one or more points with RMB, you can then move (G), scale (S) or rotate (R) them. You can also select and edit handle points that allow to finer control the interpolation curve. you can also set points values numerically by hitting N then entering a value in the "Vertex X" and "Vertex Y" fields inside the Transform properties panel.

Adding Curve Points

Mode: All Modes

Hotkey: CtrlLMB / â § ShiftD

Menu: Ipo Curve Editor → Point → Duplicate

Alternatively to adding curve points via inserting keyframes (eg. in the 3D View or buttons window), curve points can also be added manually in the Ipo Curve Editor by using CtrlLMB. The point will be added to the active curve (the one with its color button next to its channel name outlined in black even one that doesn't have any curve point already set), at the location of the mouse pointer.

Points are added according to the following rules:

- There is no Ipo datablock yet (in this window) *and* one *channel* is selected: a new Ipo datablock is created along with the first Ipo Curve with one vertex placed where the mouse was clicked.
- There is already an Ipo datablock, and a *channel* is selected without an Ipo Curve: a new Ipo Curve with one vertex is added.
- There is already an Ipo datablock, and a *channel* is selected with an existing Ipo Curve: A new point is added to the selected Ipo Curve linked with the nearest points before and after it. If multiple curves are selected, and/or you are in Edit mode, the points are added to the *active* curve (the one with the button next to its channel's name outlined with black).

Selected curve points can also be duplicated. Note that this will not duplicate the curve segments, but the points themselves, connected to the original curve.

Hints

Repeated CtrlLMB clicking along a curve can be a good way to add randomness and unpredictability to an animation

Point Handle Types

Mode: All Modes

Hotkey: H, â § ShiftH, AltH, V

Menu: Ipo Curve Editor → Point → Handle Type

Similarly to 3D Curve objects, Ipo curve points possess two handle points by each side of it. Those handle points can have a variety of types, that can be used to specifically control the interpolation of segments between points.

Options

Free Handle H (black)
: The point handles are unlinked and can be manipulated in any way. (H toggles between Free and Aligned)

Aligned Handle H
: The point handles are linked in a straight line. (H toggles between Free and Aligned)

Vector V
: Both ends of a handle always point to the previous or next handle

Auto â § ShiftH (yellow)
: This handle has a completely automatic length and direction, based on the points' proximity to each other

As soon as handles are rotated, by moving one of the point's handles, the type is changed automatically:

- Auto Handle becomes Aligned.
- Vector Handle becomes Free.

Ipo Keys

Mode: All Modes

Hotkey: K

Menu: Ipo Curve Editor → View → Show Keys

As easy as it is to work with animation curves, some people find it even easier to manipulate the keys. Changing the Ipo Editor to keys display makes two things happen:

- The Ipo Curve Editor now draws vertical lines through all the points of all the visible curves (curves are now shown in black). Points with the same 'frame' value are linked through the vertical lines. The vertical lines (the "Ipo Keys") can be selected, moved or duplicated, just like the points. You can translate the keys only horizontally.
- The object is not only shown in the 3D View in its current position but 'ghost' objects are also shown at all the key positions. On some video displays, you may have to press K in the 3D View window. In addition to now being able to visualize the key positions of the object, you can also modify them *in* the 3D View. For example, you can move the selected Ipo Keys.

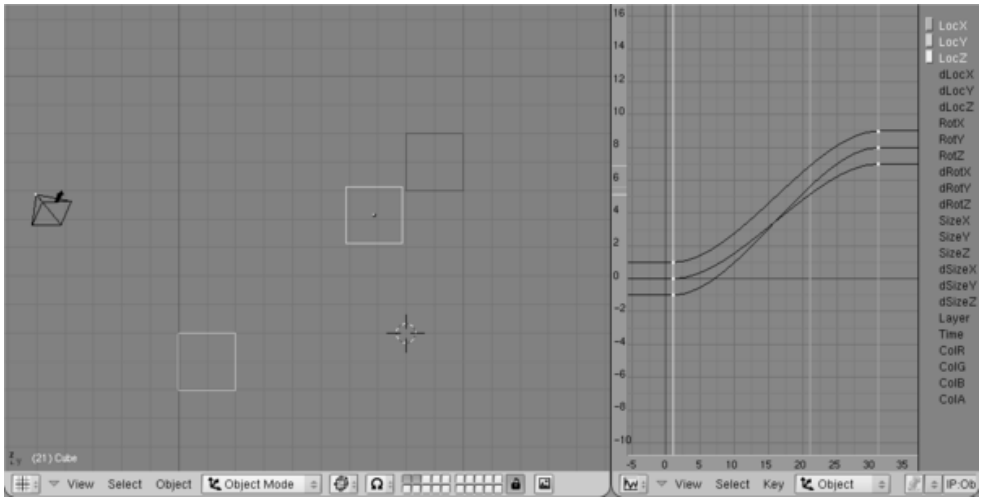
Onion Skinning and Ghosting

It is possible to see the keyed positions of an object in the 3D View. With the IPO editor in Keyframe mode, press K in the 3D View window. The location of the object will be drawn as ghostly outlines. The ghost location of the selected key will be drawn in a highlight color. Using this feature, you can physcially see the location, or path, that your object will take in 3D space.

Hints

- Use keyframe mode to change where an object is in time. For example if you currently have an object in a certain position in frame 100, but find that it gets there too early, go into key mode, select that key at frame 100, and Ggrab it and move it to the right. Click to drop it into position.
- Holding Ctrl down while dragging keyframes makes them snap to even frames.
- The I "Insert Key" always affects *all* selected objects. The IPOKeys for multiple objects can also be transformed simultaneously in the 3D View. Use the â § ShiftK command: Timeline â Show and Select Keyframes to transform complete animations of a group of objects all at once.
- Use Page up â and Page down â commands to select subsequent keys in the 3D View.
- You can create Ipo Keys with each arrangement of channels. By consciously *excluding* certain channels, you can force a situation in which changes to key positions in the 3DWindow can only be made to the values specified by the visible channels. For example, with only the channel LocX selected, the keys can only be moved in the X direction.
- Each Ipo Key consists of the vertices that have *exactly* the same frame value. If vertices are moved manually, this can result in large numbers of keys, each having only one curve. In this case, use the J ("Join") command to combine selected IPOKeys. It is also possible to assign selected IPOKeys vertices for *all* the visible curves: use I in the IPOWindow and choose "Selected keys".
- The DrawKey option and the IPOKey mode can be switched on and off independently. Use the button EditButtons->DrawKey to switch off this option for this object. You can switch IPOKey mode on and off yourself with K in the IPOWindow. Only K in the 3DWindow turns on/off both the DrawKey and IPOKey mode.

Examples



The IPOKey mode

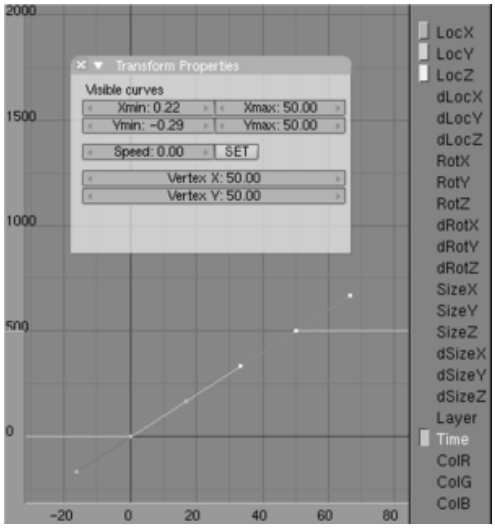
Special Notes on the Time Curve

Mode: All Modes

With the Time curve you can manipulate the animation time of objects without changing the animation or the other Ipos. In fact, it changes the mapping of animation time to global animation time (*Linear time IPO*). The Time curve is a channel in the Object Ipo.

In frames where the slope of the Time curve is positive, your object will advance in its animation. The speed depends on the value of the slope. A slope bigger than 1 will animate faster than the base animation. A slope smaller than 1 will animate slower. A slope of 1 means no change in the animation, negative power slopes allow you to reverse the animation. The Time curve is especially interesting for particle systems, allowing you to "freeze" the particles or to animate particles absorbed by an object instead of emitted. Other possibilities are to make a time lapse or slow motion animation.

Examples



Linear time IPO
To grasp this concept, make a simple keyframe-animation of a moving object, from a position to another in, say, 50 frames. Then select the Time channel and create a TimeIpo in the IpoWindow going from point (1,1) to point (50,50). It is easy to set the start and end point of an IPO by using N and entering the values numerically.

Multiple Time IPOs
You need to copy the TimeIpo for every animation system to get a full slow motion. But by stopping only some animations, and continue to animate, for example, the camera you can achieve some very nice effects (like those used in the movie "The Matrix")

Blender really excels at animation, and the Timeline window is like having your own VCR controls at your fingertips. The timeline window, identified by a clock icon, is by default in the SR:4-Sequence screen layout just above the button window.

Timeline Window

The X-axis of the window is the timeline of your animation in frames. Frames are the fundamental unit of measure of time in Blender. LMB -clicking anywhere in the window brings Blender and all the other windows to that frame in time. Moving the mouse while holding MMB or AltLMB pans the timeline to the left or right. Holding CtrlMMB scales the display (frame range) in or out by moving the mouse left or right, just like in the 3D window.

Any Ipos (keyframes) for the currently selected object appear as vertical yellow lines at the frame they occur.

The exact frame is shown in the frame counter in the Timeline header, as well as the Buttons header. To the right of the frame counter are your standard VCR controls to rewind, skip back, play or pause, skip forward and skip to the end, and record. The start and end are set using the fields to the left of the frame counter, and synchronize automatically with the start and end frames in the Render buttons underneath the Anim button.

The menu consists of choices to View, Frame or Playback

View Menu

This menu controls what you view and how you view it:

- maximize display is a standard window feature to make this window full-screen.
- If Lock Time is enabled, all other windows are automatically synchronized to this window as the master clock.
- View All automatically zooms the display to show all events (Ipo keyframes) in range, so you don't have to pan.
- Center view centers the display on the current frame.
- Jump command do the same as the buttons in the header.
- Based on the frames per second you have set in the Render buttons, toggling the Show Seconds changes the display from frame units to seconds and fractions thereof. Very often storyboards are laid out in seconds; choosing this display unit makes things a little easier than doing all that multiplying in your head.
- Play back plays the animation from start to end. It keeps playing until you click the pause button (the two vertical lines).

Frame Menu

The Frame menu does something using the frame of animation:

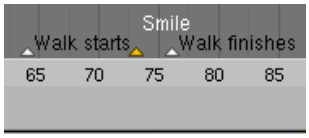
- Set as... sets the current frame selected as the start or end frame as your range.
- Setting a Marker creates an orange pointer. Controlling the name of the marker by CtrlM allows you to give a meaningful name to what is happening at that moment in time, such as "Boy loses Girl" or "Soda can spritzes". Using these Markers allows you to annotate the storyboard and animation sections. Of course, you can Grab, delete and duplicate markers.

Playback Menu

Playback controls how the animation is played back, and where.

- Set Frames/Sec allows you to change the speed of your animation.
- Windows enable the types of windows to be updated with each frame during playback. Otherwise, to conserve CPU power and allow smooth playback, some window updates can be disabled.

Markers



Markers: small but useful
Markers are used to denote frames at which something significant happens - it could be that a character's animation starts, the camera changes position, or a door opens, for example. Markers can be given names to make them more meaningful at a quick glance. When you create a marker in the timeline window, it also appears in the IPO, Action, NLA, and Audio windows, and vice versa, allowing you to see where the significant points of your animation lie, whichever window you are in.

Creating Markers

Mode: All

Hotkey: M

Menu: Timeline â Frame â Add Marker

The simplest way to add a marker is to move to the frame where you would like the marker to appear and press M.

Alternatively, you can either press AltA or the Play Timeline button to make the animation play, and then hit M} at the appropriate points. Once you stop the playback, the markers will appear. This can be especially useful for marking the beats in some music.

Once you have made a marker, use RMB to select it, and CtrlRMB to select multiple markers.

Naming Markers

Mode: All

Hotkey: CtrlM

Menu: Timeline â Frame â Name Marker

Having dozens of markers scattered throughout your timeline won't help you much unless you know what they stand for. You can name a marker by selecting it, pressing CtrlM, typing the name, and pressing the Okay button.

Moving Markers

Mode: All

Hotkey: G

Menu: Timeline â Frame â Grab/Move Marker

Once you have one or more markers selected, press G to move them, and confirm the move with LMB or â µ Enter. If you hold Ctrl while moving them, the markers will move in steps corresponding to one second - so if you have 25 fps set, the markers will move in 25-frame steps.

Deleting Markers

Mode: All

Hotkey: X

Menu: Timeline â Frame â Delete Marker

To delete the selected marker(s) simply press X and confirm with LMB.

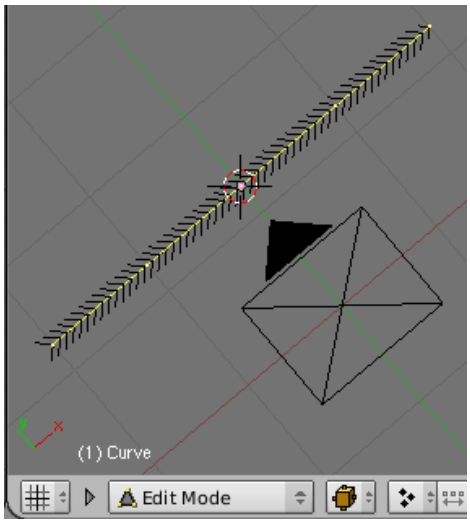
A different way to have Objects move in the space is to constrain them to follow a given path. Sometimes objects need to follow a specific path through a minefield, or it is just too hard to animate a special kind of movement with the keyframe method. For example, a planet following its way around the Sun - animating that with keyframes is virtually impossible. In these cases, a special type of curve object called a *Path* can be used as the path for another object to follow.
There are two steps in animating an object to follow a path:

1. Create the Path
2. Tell the object to follow that Path

Creating a Path

Mode: 3D View Object Mode

Hotkey: Space → Add → Curve → Path



Path in Edit Mode
A Path is a special type of Curve object. Blender has a specific Path curve object, added as indicated in the reference box above, but any curve can be used as a path. So, for example, if you used a curve to show the groove in the toy track of a electric race car, you can use that same curve as the path for the race car itself. The two act the same, however there are some automatic features that make using a Path easier than a generic Curve.

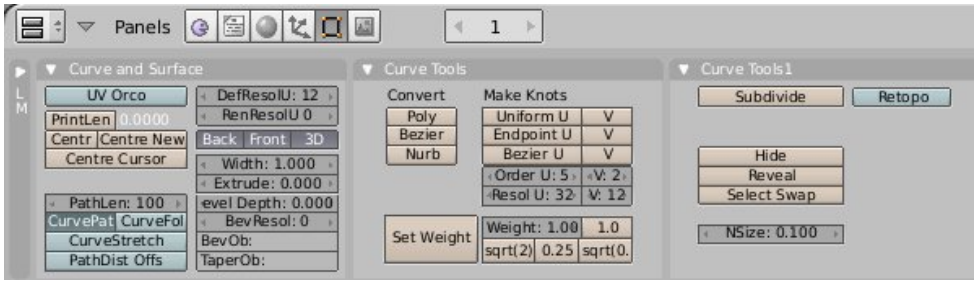
The Path is made up of control points, initially 5, and has a linear direction shown by the the arrows in Edit mode. You can close a path by pressing C to make the path cyclic. In the example to the right, the path is going in the +X direction in a straight line (left to right).

Speed IPO: When you add a Path, a Speed IPO curve is also automatically added, shaped such that any object traversing the path smoothly accelerates at the start and then decelerates smoothly to a stop at the end of the path over 100 frames. You can view this Speed IPO curve in the IPO window, Path IPO type. This Speed curve overrides the Pathlen value. You can edit this IPO curve to change the speed and location of the object over whatever time period you wish, using the standard IPO editing features of Blender. The Y value represents the "percent complete" of the object along the path; 0.0 is the start of the path, and a value of 1.0 is the end. The X value is the frame of the animation. Thus, a sine-wave-shaped (sinusoidal, for you geometry geeks) Speed curve will make any objects that follow the path to move back and forth along the path in a very hypnotic fashion.

Like all objects, Paths are named in the Links and Materials panel or in the Properties panel, by default "Curve". This becomes important when you constrain an object to the path, because you have to enter the path's name in the Ob: field in the constraining panel. The older way to make an object follow a generic curve is to parent the object to the curve. This legacy technique is described below. The rest of this paragraph discusses using the Path object and Object Constraint method.

Select the object which you want to follow the path. In the Object (F7) context Object buttons, click Add Constraint->Follow Path. In the Target OB: field, enter the name of the path. The Constraint name button should turn from red to grey, and the center of the object will jump to the start of the path. Click here for more information on the constraint (see also its dedicated page in this manual).

Path Buttons Panels



The Curve/Path Editing Panels
The path has all of the normal curve tools available. Focusing on useful Path features, working left to right through the panels:

Curve and Surface Panel

> **CurvePath** designates the curve as a path that can be followed (on by default). Objects can be translated (location) by the path, and if **CurveFollow** is enabled, also rotate to point along the path, like a train on a track. Similar to a Lattice, **CurveStretch** makes the object stretch, like a rubber band, along the path. **Pathlen** defines the default number of frames over which the object will traverse the path. A Speed curve provides much more control and, if present, overrides the Pathlen value.

Curve Tools Panel

> You can change the type of curve (Poly, Bezier, Nurbs) which changes the way the control points influence the path. Each Control point has a weight, or influence, on the speed of the object along that part of the path. The preset buttons save you from typing numbers in the numeric field. The square root buttons calculate smooth acceleration/deceleration using the square root of 2 as a speed factor between points. Be sure to click Set Weight to lock in your change to that point.

Curve Tools 1 Panel

> The Subdivide button is useful for inserting more control points between selected control points.

Changing the Path

Mode: 3D View Edit Mode

Hotkey: G to Grab, E to Extrude, C to toggle Cyclic

Menu: Curve → Extrude; Curve → Toggle Cyclic, Curve → Segments → Switch Direction

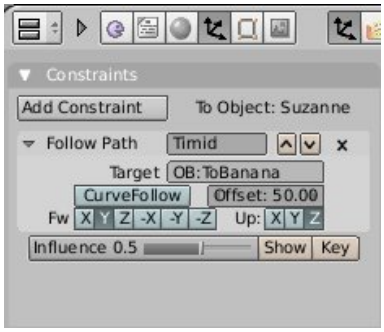
All the normal vertice editing controls are available, mostly Grab. Special Curve tools (like Cyclic, which make a closed path) are also available. Of special note:

Extrude
: used to extend either the start or the end of the path. With the start or end point selected, either Extrude or CtrlLMB click to add points to that end of the path.

Switch Direction
: swaps the path start and end points

Anim settings Panel
: the TimeOffset numeric at the bottom of this panel specifies when in time this path starts influencing its objects

Following a Path



Path Constraint on Object
With your object selected, add a Follow Path constraint to the object. Enter the name of the Path, and specify the orientation. Paths have a local XYZ orientation, and the buttons align your object's local orientation to the path. In this example to the right, our favorite monkey has a "Timid" constraint to follow the path named "ToBanana". The path only influences her motion by about 50%, so she does not follow the whole path, but is only displaced by half the path length. She starts along the path to the banana at frame offset 50. The actual start frame is this offset plus the Path's animation setting for TimeOffset.

When you add a constraint, the object moves relative to the path, but does not have to physically follow that path in the same 3D space. It can be physically apart (beside, away from) the actual curve. This allows you to have a generic path, like "Falling in the wind" and then model many Leaves ("Leaf.001", "Leaf.002", etc.). Each leaf is moved away from the path and thus follows the path, but in its own space. Any Ipo Curves for that object are still in effect as it moves along the path, allowing you to let the object swivel left or right as they walk along the path, like a cartoon bird spinning and sliding on the ice.

Speed/Progress via Time Curve

You control the amount of time (actually the object's percentage progress) by RMB selecting the curve and, in the IPO window, LMB selecting the Time channel. Then CtrlLMB click to add interpolation points at

various frames, with the Y value between 0 and 100 reflecting the progress or distance along the path the object should be at that time (frame). If the curve is increasing, motion is forward; if the curve is decreasing, motion is backward. Thus, a simple sine wave (extended cyclic) curve will make an object swing back and forth along the path forever.



Paths can follow other Paths:

Just add the FollowPath constraint to a path. For example, a standard "out of the ballpark home run" baseball trajectory can be offset by the wind path.

Legacy Path Method

Blender 2.37-+

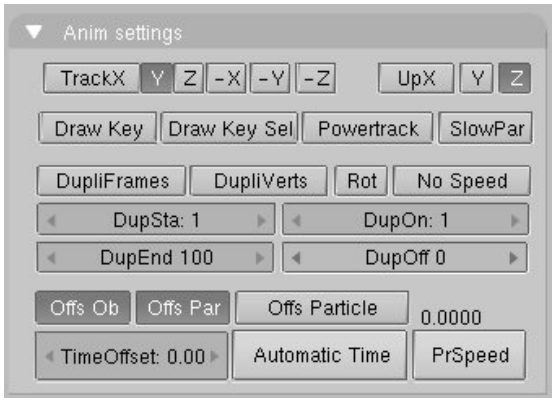
Editor's Note

The following text was in place prior to Version 2.43. It is preserved here for users of older versions, and for working with old blend files that use the old method of path animation. The spaceship example is still relevant and helpful in explaining/providing more detail.

As for tracking, there are two Path animation methods, the old, pre 2.30 method described here and the new method, which actually defines a constraint, which will be described in *character_constraints*.

When parenting an Object to a Curve, you will be asked to choose a Normal Parent or a Follow Path option. The Former is what you need for conventional Path animation, but other actions needs to be taken later. The second option creates a "Follow Path" Constraint, and it is all you need to do. If, when following a path, the dashed line between the object and path goes to the wrong end of the path simply clear the parent, select the path and go into edit mode. Select all the control points and select, in the Curve Menu, segments and select switch direction, then re-parent the object to the path. Any kind of curve can become a path by setting the option CurvePath Toggle Button in the Animation Buttons window (F7) to ON (*The Action Window with Path Buttons*.).

When a Curve has children objects it can be turned to a Path by selecting it, going int0 the Editing Context (F9) and activating the CurvePath Toggle button in the Curve and Surface Panel. Child objects of the Curve will now move along the specified path. It is a good idea to set the Curve to 3D via the 3D Toggle Button of the Curve Edit Buttons so that the paths can be freely modeled. Otherwise, in the ADD menu under Curve->Path, there is a primitive with the correct settings already there. This is a 5th order NURBS spline, which can be used to create very fluid, continuous movements. Normally a Path is 100 frames long and it is followed in 100 frames by children. You can make it longer or shorter by varying the PathLength: Num Button. The *speed* along a path is determined with an appropriate curve in the IPO Window. To see it, in the IPO Window Header button you must select the Curve type for the IPO block. A single channel, Speed is there. The complete path runs in the IPO Window between the vertical values 0.0 and 1.0. Drawing a curve between these values links the time to the position on the path. Backward and pulsing movements are possible with this. For most paths, an IPO Curve must run *exactly* between the Y-values 0.0 and 1.0. To achieve this, use the Number menu (N) in the IPO Window. If the IPO Curve is deleted, the value of PathLen determines the duration of the path. A linear movement is defined in this case. The Speed IPO is a finer way of controlling Path length. The path is long 1 for time IPO, and if the time IPO goes from 0 to 1 in 200 frames then the path is 200 frames long. Using the option CurveFollow, in the Curve and Surface Panel, a rotation is also given to the Child objects of the path, so that they permanently point in the direction of the path.



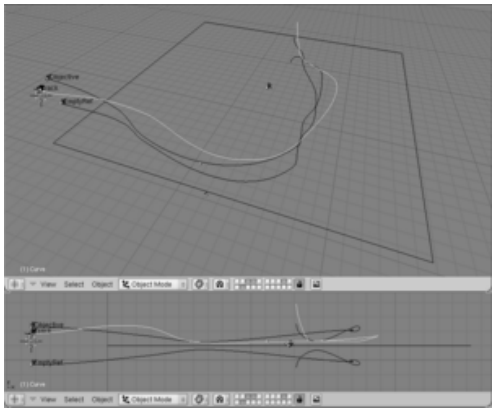
Tracking Buttons

Use the "tracking" buttons in the Anim settings Panel of the Object (F7) context to specify the effect of the rotation (*Tracking Buttons*) as you would do for Tracking:

- TrackX, Y, Z, -X, -Y, -Z This specifies the direction axis, i.e. the axis that is placed on the path.
- UpX, UpY, UpZ Specifies which axis must point 'upwards', in the direction of the (local) positive Z axis. If the Track and

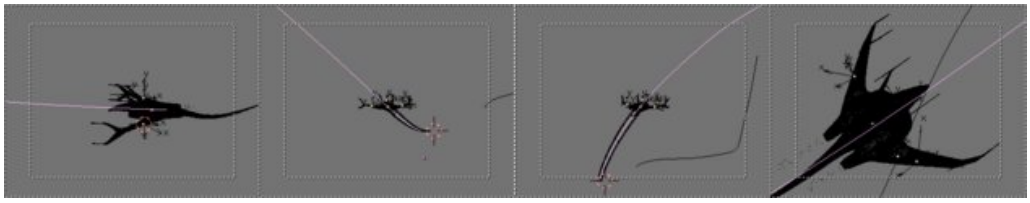
the Up axis coincides, it is deactivated.

Note
Curve paths have the same problem of Beveled curves for what concern the definition of the "Up" direction.



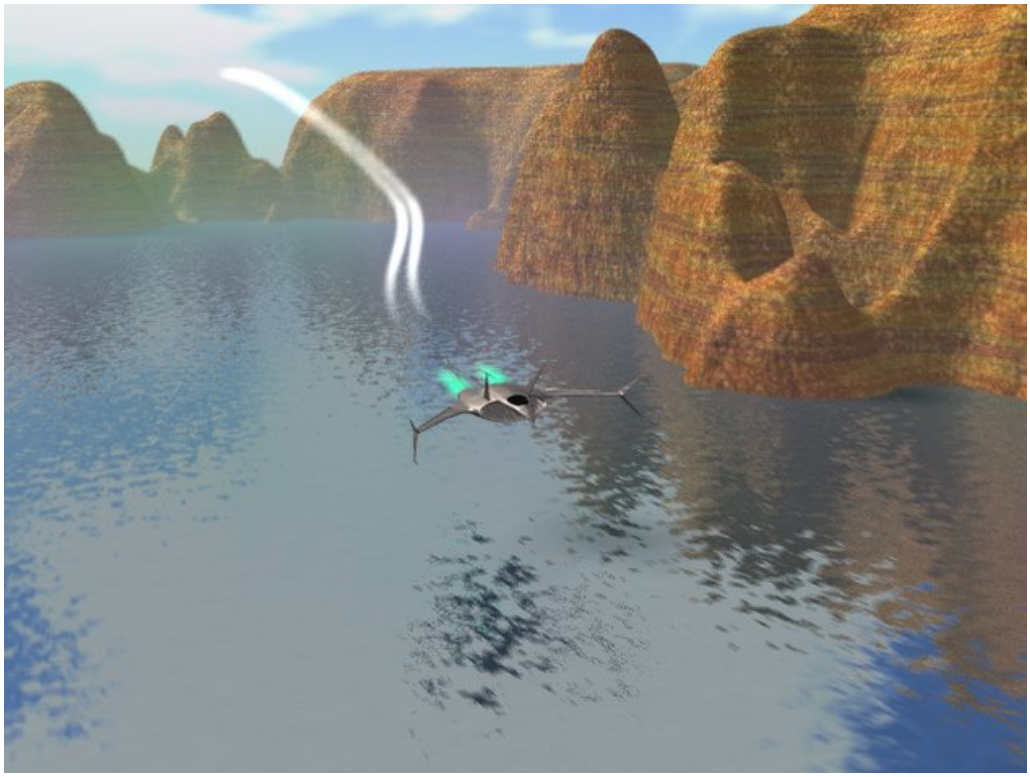
Complex path animation (Click for larger image)
To visualize these rotations precisely, we must make it possible for a Child to have its own rotations. Erase the Child's rotation with AltR. Also erase the "Parent Inverse": AltP. The best method is to 'parent' an *unrotated* Child to the path with the command Ctrlâ § ShiftP: "Make parent without inverse". Now the Child jumps directly to the path and the Child points in the right direction. 3D paths also get an extra value for each vertex: the 'tilt'. This can be used to specify an axis rotation. Use T in EditMode to change the tilt of selected vertices in EditMode, e.g. to have a Child move around as if it were on a roller coaster. *Complex path animation* shows a complex application. We want to make a fighter dive into a canyon, fly next to the water and then rise again, all this by following it with our camera and, possibly, having reflection in the water! To do this we will need three paths. Path 1 has a fighter parented to it, the fighter will fly following it.

The fighter has an Empty named 'Track' Parented to it in a strategic position. A camera is then parented to another curve, Path 2, and follows it, tracking the 'Track' Empty. The Fighter has a constant Speed IPO, the camera has not. It goes faster, then slower, always tracking the Empty, and hence the fighter, so we will have very fluid movements of the camera from Fighter side, to Fighter front, other side, back, etc. (*Some frames, the camera fluidly tracking the fighter.*)



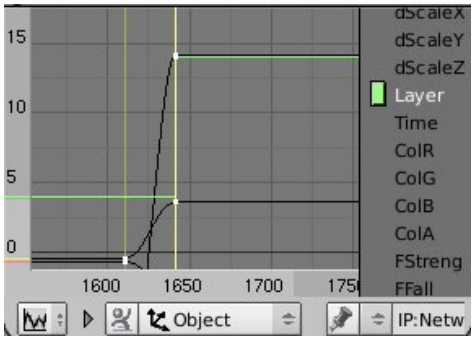
Some frames, the camera fluidly tracking the fighter. (Click for larger image)
Since we want our fighter to fly over a river, we need to set up an Env Map for the water surface to obtain reflections. But the Empty used for the calculations must be always in specular position with respect to the camera... and the camera is moving along a path! Path 3 is hence created by mirroring path 2 with respect to the water plane, by duplicating it, and using M in Edit Mode with respect to the cursor, once the cursor is on the plane. The Empty for the Env Map calculation is then parented to this new path, and the Time IPO of Path 2 is copied to Path 3. *A frame of the final animation.* shows a rendered frame. Some particle systems were

used for trails. The scene presents many subtle tricks, as particles for the jet streams, fog, a sky sphere encircling the scene and so on.



A frame of the final animation.
Many times you want an object or actor to wait off stage until it is their time to come on and do their thing. In Blender, you commonly render only a few of the 20 available layers. Further, you can set some lamps to illuminate only objects on their layer. By changing an object's layer during the animation, it is possible to animate an object's layer "membership" and thus it's visibility in the scene, and its lighting.

Further, by putting the object on a layer that is not selected in your 3D view, it will not clutter up your workspace. During the frames of the animation when it is on a visible layer, you will see it and be able to work with it.



Changing from Layer 4 to 14
To animate an object's layer membership, go to the frame when you want the object to appear. Ensure the object is on a visible layer (Move Layer command will tell you what layer it is on). Insert an Ipo Key, and choose Layer (the menu option located below the Location and Rotation selections). In your Ipo Window, you will see the Layer channel now has a value corresponding to the Layer. Back up one frame, Move the object to an unselected (invisible) layer. Now, here's a slight glitch; you cannot add a Layer Key because the object is on an unselected layer. So, temporarily enable that layer by shift-selecting it from the 20 Layer buttons, Insert another Layer key, and then de-select the undesired layer.

Select All Layers
To activate all Layers, press ` (accent grave .. top left of US keyboards)

Scrub or play through the animation and you should see your object appearing at the indicated frame. You can edit this Ipo control point, usually in Keyframe mode, to change when the shift occurs.

Tips

You will note that the Layer Ipo "curve" is a discontinuous jump. There is no way to have an object gradually change layers. To make an object fade into view, you have to animate its Material Alpha setting to fade in from 0.0 to 1.0.

There are two ways to check which layer an object is on :

- Press M to display the layers menu, the object's layer(s) buttons will be enabled. Press ESC or OK to dismiss the menu

- Display the Buttons Window / Object Buttons F7. Under the Draw panel, the Layers buttons displays the object's active layer(s). The active layer(s) can be changed from that panel as well, instead of using the floating layer panel M

Animating an Object's location, rotation, scale are trivial animations. Animating its Material settings is a bit more conceptually challenging. So let's take it to the next level, shall we? With Blender, you can change, reshape, and deform your objects over time!

There are many ways of achieving this actually, with increasing levels of complexity. However, each level builds on one another and as usual, there are many ways to do it. Depending on how fancy you want to get and what you want to do, one approach may be preferable over another.

- Hooks use an Empty (or other object, but most commonly an empty) to pull at selected vertices of your mesh or curve or lattice or NURBS surface. You animate the location of the Hook to animate the deformation of your mesh. Hooks are modifiers and are discussed in the Modifer section of this manual.
- Relative Shape Keys change the location of some of the vertices of a mesh. Relative shape keys can be combined, sometimes called BlendShapes, because different shapes are blended together. For example, a face mesh can have one shape key to smile, and another to blink. Those two action can be combined so that it smiles and blinks at the same time.
- Absolute Shape Keys allow you to deform and change the *entire* mesh, "morphing" from one position to the next. (You can even add or subtract vertices over time.)
- Lattice Animation uses a lattice as a cage that uniformly squeezes or expands a mesh, like the event horizon in a black hole, or as a comical cartoon car screeching to a stop. You define shape keys for the Lattice, which in turn affects the target mesh as it moves through the lattice.

The latter method, Lattice Animation, brings up an interesting feature in Blender, in that one object can deform another. This feature has application especially to Bevel and Taper objects. These are curves that affect other objects, like meshes or text objects. Changing the shape of these curves over time, using shape keys, will cause other objects (the mesh, text, etc) to change shape as well.

Shape keys store different shapes of the same mesh. In other 3d applications they are called 'morph targets' or 'blend shapes' or even 'vertex keys' in older versions of Blender.

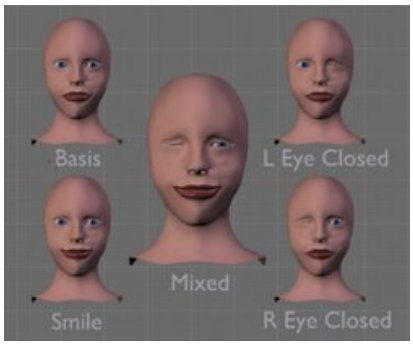
Blender offers both **relative** shape keys and **absolute** shape keys:

- *Relative* shape keys, discussed on this page, record each shapes **difference** (a.k.a. "delta") relative to a base shape, and allow one or more of these differences, as represented by their respective keys, to be combined.
- *Absolute* shape keys, as described here, record directly each different shape. Absolute shape keys apply to the *entire* mesh and are used to "morph" the mesh between one shape and the next.

Mode: Object Mode

Panel: Editing Context → Shapes

Description



An Example of Mixed Shape Keys
Relative Shape keys can be blended on a percent basis with other shape keys to achieve the desired effect. Take for example a face. The user may model a face with a neutral expression and have shape keys for a smile and a frown. These shape keys could then be combined in the Action Editor and your talking face could have some serious emotional problems.

Use shape keys for animating faces or when you want the shape of any mesh to be able to change over time; for example, a ball squishing, or a soda can being crushed. When animating faces, suggested shape keys are:

- Mouth forming each sound shape
- Lips forming a smile and frown
- Eyes: Blinking, showing basic emotions, and even winking.

Notice that a relatve shape key simply defines a "final position" for some group of vertices: it does *not* say how fast the shapes are transitioned or in what order. You can make these keys in any order or speed you wish, and control the degree of influence (to make a "slight" smile, for example), using the familiar mechanism of IPOs.

Workflow

The overall process for using shape keys in Blender is:

1. . Define your basis shape.
2. . Identify the list of primitive morphs that you might want
3. . Create Vertex Groups for sets of vertices to participate in each morph family

4. . Define your target shape(s), one for each morph
5. . Using the IPO curves, define when you want the target shape to influence the basis shape.

Target shapes can be created at any frame; they are time-independent. They exert their influence according to the shape IPO curve. Each target shape has a channel in the IPO window. When a target shape's IPO curve is at 0.0 y value, it does not deform the mesh. As its influence increases, it deforms the mesh at that frame. Since the IPO channel is a transition (a smooth curve or a linear line or whatever options you choose for the IPO curve) you can control how fast the mesh morphs to the new shape.

Multiple shapes can simultaneously morph the basis mesh, each of them exerting their influence through their IPO curve, since each of them have their own channel.

If you want to change the basis shape, you will have to create a new shape, and using a script, copy the shape keys over to the new shape, adjusting them as you desire.

A morph family is a group of shapes related to an area of the mesh. For example, take the face. You have the eyebrow area (possible shapes are Surprise, Furrow, eyebrow lift-left, and eyebrow lift right), the eyelids (Blink), the mouth (smile, frown), nose (flare, wiggle) and tongue (ay, th, i, stickout, tunnel). The jaw open action may be controlled by a bone. In this case you would a vertex group for the eyebrow, another for the tongue, etc.

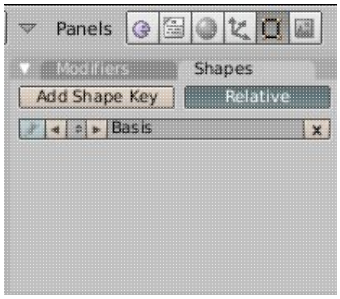
Creating Relative Shape Keys

Mode: Object Mode with Mesh, Curve or Lattice object selected

Panel: Editing Context → Shapes

Hotkey: I

Menu: Object → Insert Keyframe



Shape keys are all stored relative to a basis mesh. You can enable shape keys and insert the basis key by selecting the mesh object and pressing the I then choosing Mesh. Alternately, this step can be done with the Add Shape Key button in the Shapes panel. This panel is located as the second tab in the Editing F9 buttons in the Buttons window. In this panel, be sure that the Relative button is enabled.

The original shape key is shown in the Ipo Window, (Ipo Type: Shape), for relative keys as an ellipses just under the pushpin. The original absolute shape key is named "Basis". This difference alerts you to whether you are working with Relative or Absolute shape keys.

Once the basis shape key has been created, you can create additional shape keys in the same way. With the mesh selected, and in its deformed state, either press I and then choosing Mesh, or click the Add Shape Key button in the Shapes panel. A new Shape Key will be created. If this is the first shape created after the basis shape, it will be called Key 1. â § ShiftLMB into the name field in the panel to change the name to something meaningful.

An important thing to remember is that when creating a new shape, the new shape will be based off the current selected shape **WHEN YOU ENTER EDIT MODE**. When creating shapes from the basis shape, the new shape will be a copy of that shape, if the user creates a new shape while another altered shape is selected, the new shape will be a copy of that altered shape. This allows you to further edit selected shapes; simply leave edit mode to lock in that shape.

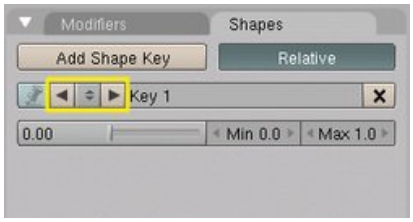
Fun Fact
Old school Blender users sometimes refer to shape keys as Relative Vertex keys, or RVKs for short.

Editing Shape Keys

Mode: Object Mode

Panel: Editing Context → Shapes

Various aspects of each shape key can be edited by activating it in the Shapes panel.



Location of the Shape Key selection tools in the Shapes panel

Arrows and the Menu
: Choose the active shape key. The arrow buttons cycle through shapes, and the menu allows you to choose a specific shape directly. Note that when you first choose a key, it is visualized at full strength in the 3D View regardless of any key frames that have been set, so you can see exactly what shape you're activating. Changing the current frame will reset the 3D View to show the results of the shape key blending again.

 Once the desired key is selected, the shape can be altered by entering editmode â Tab and manipulating the vertices. **Note:** Adding or deleting a vertex once shape keys have been added is very tricky. The changes are propagated to the other shapes based on their position in the current shape, this can have drastic effects on other shapes.

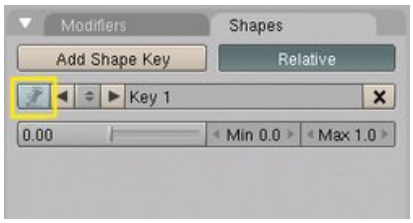
Current Shape Key Name (text field)
: Rename a shape key from its automatic name by â § ShiftLMB in the name field and typing in a new name to change it. It's very important to keep your shapes organized and well named, so you can remember which one you're editing in the Action or Ipo editor. It's very easy to end up with a huge amount of shape keys when you start animating faces and you don't want to get your work confused.

Propagating Changes to other Shapes
: If you find an error in your basis shape *after* you have started defining other shapes, you can make your changes and, with those vertices selected, use the specials menu W -> Propagate to All Shapes" in Edit mode. Those vertex positions will replace the corresponding vertex positions in the other Shape keys.

Propagating Changes from another Shape

If you find an error in your basis or other shape *after* you have started defining other shapes, with those vertices selected, use the specials menu W -> Blend from Shape and choose which shape key to use as a reference for repositioning the vertex. You can then even choose to what extent to apply that reference point to the existing vertex location. Think of it as a way to average the position of vertices between shapes.

Pinning



Location of the Pin Icon in the Shapes panel

When a mesh object does not have a pinned key, it can show multiple keys at once. To prevent this, click on the pin icon.

The current mesh object is now locked at that shape key, and will only show that key. This feature is useful when multiple keys are affecting a mesh object and you would like to see the effect of one specific key.

One other use for pinning is to create a shape gallery.

Since Blender allows for the creation of Linked Duplicates which share underlying data (e.g. mesh and shape key data) multiple copies of a mesh can be created using AltD.

Once a duplicate has been created, move the duplicate to a different area of the screen. Pin the duplicate to a shape key by clicking the pin icon in the Shapes panel. Finally, reselect the 'original' mesh. If this is done for each key, a shape gallery emerges (which also helps in editing keys).

Hints

Pinning Shape Keys can also help to speed up interaction with a mesh (especially when animating it with armatures). If a shape is pinned, Blender doesn't have to calculate the myriad interactions and blends of different keys in real time, which speeds up operation. This is particularly useful when using armature-driven shape keys to correct deformations, since it's usually not something critical that you need to see when animating. Remember to un-pin the shapes when you're done and ready to render though!

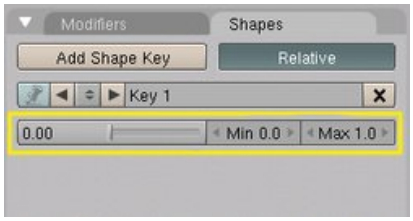
Blending Shape Keys

Mode: Object Mode

Panel: Editing Context → Shapes

Shape Keys can be mixed in several different ways. These methods have the same net effect of creating IPO data, but are different in interface.

Blending Using the Shapes Panel



Shape Panel and Shape Keys
In the Shapes panel when a key is not pinned, it has an additional row of buttons; value, min and max.

Key Value (slider)
: Current value of that key at the current frame. Adjusting this slider will insert or adjust a control point in the shape key control curve at the current frame. If I put my "blink" shape key at 50% instead of 100%, I'll have half-open eyes.

Min and Max
: Adjusts the boundary values for the value slider both in the Shapes panel and the Action Editor Window. It is possible to use a maximum greater than 1.0, to push shapes to even greater extremes, but be careful, since you are exaggerating the shape key, and this can mesh up your mesh.

 It is often easiest to to mix your shape keys in the Action Editor instead of doing it from the main editing panel.

Blending Using the Action Editor



Action Editor and Shape Keys

In the Action Editor, select "ShapeKey editor" mode, expand the Sliders button and use the sliders to adjust the key values at the current frame. The min / max values set in the Shapes panel are reflected in the Action Editor window sliders for that shape.

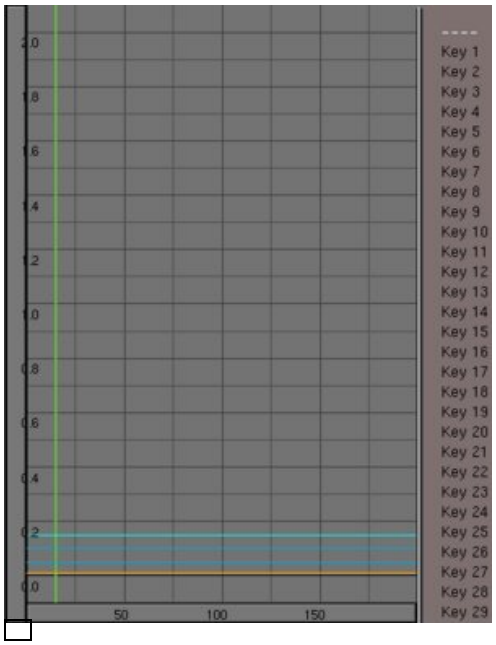
If there is not already a keyframe at the current frame (indicated by the green line), a new keyframe will be added.

Blending Using the Ipo Curve Editor

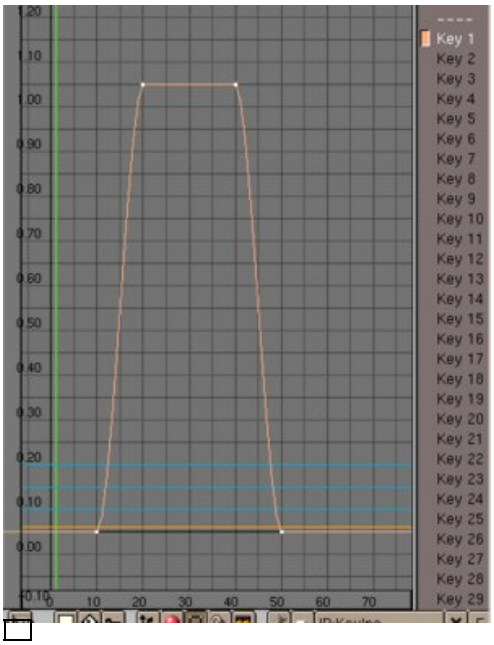
The Ipo Curves show the underlying data that is being controlled by the other mixing methods. It is in this window where the values are mapped to frames for each shape key. In order to access the Shape key Ipos, select Shape from the Ipo type popup menu in the Ipo Curve Editor Window header. Each shape key will be shown in the Ipo list on the right side and can be altered in the same manner as any Ipo curve in Blender. See the "Editing Ipo Curves and Keyframes" section for more information.

The vertical order of the Vertex Keys (the blue lines) from bottom to top determines its corresponding Ipo Curve, i.e. the lowest blue key line will be controlled by the Key1 curve, the second lowest will be controlled by the Key2 curve, and so on. No Ipo is present for the reference mesh since that is the mesh which is used if all other Keys have an Ipo of value 0 at the given frame.

The Ipo Curve of Key 1 shows a mesh undeformed up to frame 10, then from frame 10 to frame 20 Key 1 will begin to affect the deformation. From frame 20 to frame 40 Key 1 will completely overcame the reference mesh (IPO value is 1). The effect will fade out from frame 40 to frame 50.



Keys in the IPO Window



The IPO curve of Key 1

Using IPO Drivers

See the Advanced Driven Shape Keys page for information regarding how to set up an IPO Driver.

Absolute Shape Keys

Mode: Object Mode / Edit Mode

Panel: Editing Context → Shapes

Hotkey: I

Menu: Object → Insert Keyframe

Description

As previously discussed, there are two kinds of shape keys:

- *Relative* shape keys, discussed here, allow you to define and then blend individual deformations (such as a "wink").
- *Absolute* shape keys, discussed in this page, allow you to morph (only) the *entire* mesh between one or more target shapes.

Shape Keys (as opposed to Object keys, the position of the whole object center) are the specified positions of vertices *within* an Object; the actual points that define the mesh. Since this can involve thousands of vertices for an object, separate motion curves are not created for each vertex because it would overload your computer's memory. A more generic keyed position system is used instead; a single IPO Curve is used to determine how interpolation is performed and the times at which a Shape Key can be seen.

What Types of Objects can be Shaped?

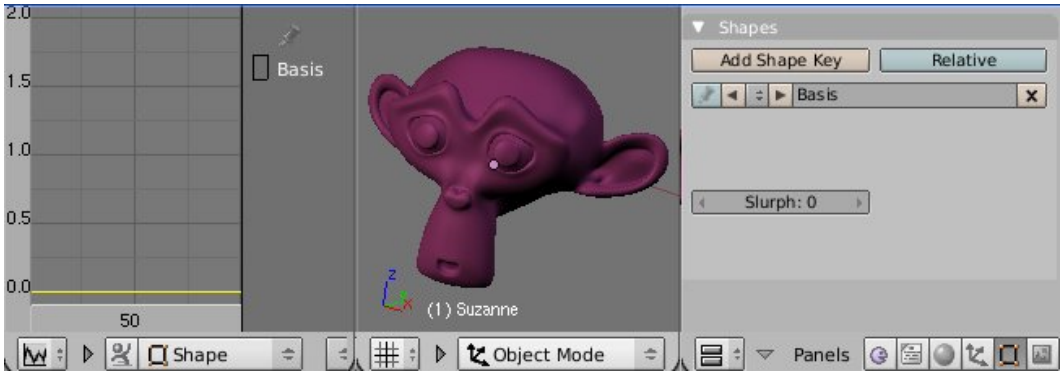
You can define Absolute Shape Keys for:

- Polygonal Meshes,
- Bezier and Nurbs Curves,
- Nurbs Surfaces,
- and Lattices.

All of these objects have vertices. The interface and use for Shape Keys is the same for all of them. In general, you define a Basis key, edit the object to a new shape, leave edit mode and insert a key for each deformation. To animate the shape shift, you then say when (in terms of Frames) you want the object to assume that shape using the Action Editor.

(For simplicity, this topic will usually just use the word "meshes.")

Creating Absolute Shape Keys

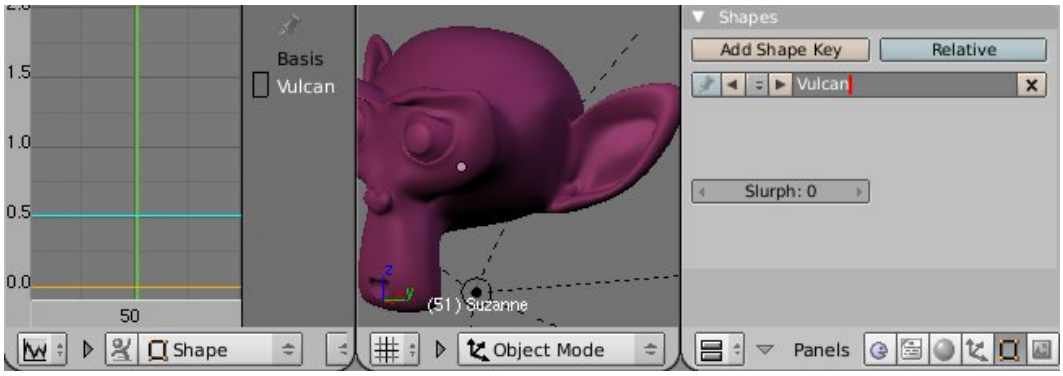


Basis Absolute Shape Key
The first absolute shape key that is created is always the *reference* or Basis key. This key is the Orange line in the Ipo Window (yellow if selected). Only if this Key is *active* can the faces and curves, or the *number* of vertices, be changed.

The subsequent shape keys are defined relative to the first. In other words, they record how a vertex has been moved relative to the Basis key position. (This doesn't affect your workflow, but it's the underlying reason why the Basis key is different from all the others. Subsequent keys are *not* defined relative to "each other." Rather, all of them are individually defined relative to the Basis key.)

A few notes about the above picture: When working with Shape keys, it is very handy to have an IPO Window open. Use the first Animation Screen layout if you like (CTRL+1). In the IPO Window, we must then specify that we want to see the Shape Keys by selecting that Ipo Type from the Ipo Window header. To do this, use the IPO type Menu Button and select Shape. Go to the 3DWindow and select the object (Suzanne, in this case) and press I. The "Insert Key" menu has several options, the last being Mesh. For other kinds of objects, this menu item will be named for that kind of object e.g. if a curve object is selected, the menu item on the Insert Key menu will be "Curve". The 3D View header menu is also context sensitive, and the menu name will change based on the type of object selected.

As soon as this has been selected, a new panel appears in the Editing panel (F9) in the Shapes panel. It has buttons to Add Shape Keys (what you just did), and a Relative button, enabled by default. To work with Absolute shape keys, disable this button. When it has been turned off (a lighter shade of green), Absolute Keys; an orange horizontal line is drawn in the Ipo Window in Shape display mode. This is the first key and thus the Basis key. The name of this shape key changes in the Ipo Window to "Basis".



Suzanne becomes Vulcan
To add successive shapes, you

- Go to the frame in the animation where you want that shape to be fully formed
- Modify the shape (in Edit mode)
- Leave edit mode
- Click the Add Shape Key button (or do the I thing again in the 3D View).

For each key you add, blue-colored lines appear in the Ipo Window (cyan when selected). They are placed above one another in a height that is 1/100 of the frame number where they were created, just to space them out (a shape created at frame 50 will appear at 0.5).

Feel free to move the shape-key lines up or down as you wish. Or, re-shape the IPO curve as you wish. Or both. Ultimately, the computer will apply the shape-key information like this:

- Each shape key (horizontal line...) defines "a target shape." (From the computer's point of view, it does so either directly ("Basis key") or indirectly ("non-Basis key"), but the result is the same. It's a shape.)
- The horizontal axis is (as usual), a frame-number...
- So, at this point, the IPO curve is either "exactly on top of" a shape-key (horizontal) line, or, more likely, "somewhere in-between" exactly two of them.
- The *ratio* between the two vertical distances, between "this point on the IPO curve" and (only) the shape-key lines directly above and directly below it, determines the shape-change now in effect. (For instance, if the point is exactly halfway between the two lines, their influences will be exactly equal.)

ShapeKey creation
Creating Shape Keys in Blender is very simple, but the fact that the system is very sensitive in terms of its configuration can cause a number of 'invisible' things to happen. The following rule must therefore be taken into consideration. As soon as a Shape Key position is inserted it is *immediately active*. All subsequent changes in the Mesh are linked to *this* Key position. It is therefore important that the Key position be added *before* editing begins.

VertexKeys can be selected in the IPO Window or in the Shapes panel using the <> buttons. We always do this *out* of Edit Mode: the 'contents' of the VertexKey are now temporarily displayed in the Mesh. We can edit the specified Key by starting Editmode.

As you render your animation over the frame range, the shape changes from one to the other. You can verify this by scrubbing over the frame range using the Timeline window or by wearing out your arrow keys.

Sharing Shape Keys

Shape Keys are part of the Mesh data, not of the Object. When cloning an object, by making both objects use the same Mesh (Editing buttons, Links and Materials panel), the associated Shape Keys are also copied. Thus, it is possible to permit multiple Objects to share the same Shape Keys in Blender by making them use the same Mesh.

Workflow for Creating Absolute Shape Keys

Once the shapes are defined, there are three methods for working with Shape Keys:

The 'performance animation' method.

- This method works entirely in EditMode, chronologically from position to position:
- Insert Key. The reference is specified.
- A few frames further: Insert Key. Edit the Mesh for the second position.
- A few frames further: Insert Key. Edit the Mesh for the third position.
- Continue the above process...

The 'editing' method.

- We first insert all of the required Keys, unless we have already created the Keys using the method described above.
- Blender is *not* in EditMode.
- Select a Key. Now start EditMode, change the Mesh and leave EditMode.
- Select a Key. Start EditMode, change the Mesh and leave EditMode.
- Continue the above process....

The 'insert' method

- Whether or not there are already Keys and whether or not we are in EditMode does not matter in this method.
- Go to the frame in which the new Key must be inserted.
- Insert Key.
- Go to a new frame, Insert Key.
- Continue the above process...

While in EditMode, the Keys cannot be switched. If you attempt to do so, a warning appears.

Shifting between Shapes

Now that you have your different shapes defined, you now want your mesh to shift between them at certain times in your animation. As you might have guess by now, there are a few ways to do it. I call them the line way, and the control way.

In the line way, in your Ipo Window, LMB click on a Shape Key; its name will turn white. Alternatively, scroll to the shape in the Shapes panel (its name won't turn white in the Ipo Window, but the box next to its name will shift to it, letting you know it is active). In the Ipo window workspace, Grab and move your mouse up or down, and drop the blue line at the frame you want.

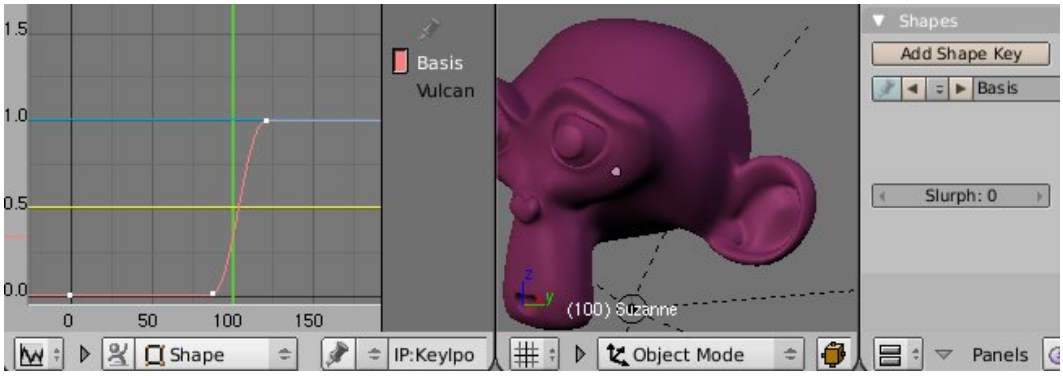
The Y-axis values range from 0.0 (0%) to 1.0 (100%), and this placement by itself represents "percent of the total time." If you are using "the line way" and simply want a linear transition between the keys, this is all that you have to do.

In the control way, you define an Ipo curve which is a sort of Time (or Speed) curve for the mesh. This curve is linked to the "`Basis`" channel (or key) â← you can define a curve for any key, but only the basis one is useful (the others have no effect). This curve maps the current frame (x axis) to a 'virtual' frame (a position in the sequence defined by the absolute shape keys) on the y axis.

Here's how it works: where the Ipo *crosses* a blue line, the corresponding key has its maximum effect and the adjacent keys have none. When the Ipo value is *in between* two keys (in between two lines), the mesh is deformed by the interpolation of *(only)* these two keys: the one "above" it and the one "below." The closer the value to a line, the higher the influence of the corresponding key, and therefore, the lesser the influence of the "other" one.

As usual, when the IPO curve has a positive slope (a "climbing" one), the animation is the same as without any curve (it might just be faster/slower, and/or have a time offset). If you give your curve a negative slope, the animation will run backwards. If you make it go up-down, the animation will go forward and backward.

So here's the effective difference between "the line way" and "the control way": "the line way" implicitly uses an Ipo that's just a linear, upward sloping, 45-degree diagonal line.



Using Control Points to Time the Shape Shift
In the example above, without any control, Suzanne would shape shift to a vulcan between frames 50 and 100 (yellow line at 0.5, blue line at 1.0). However, we selected the Basis key and added some control points (evidenced by the orange selector block). There are three control points, the important ones being at frame 80 with a value of 0, and at frame 120 with a value equal to the Vulcan key. The cursor (green line) is at frame 100. Normally, Suzanne would be a Vulcan by now, but because of the control points, she has only just started, and won't become a fully formed Vulcan until frame 120. (Because the shape between the two points is slightly curved, you'll see a bit of "ease-in ease-out" going on, too.)

To repeat: at any point in time, the IPO curve is always (by definition) "somewhere in-between" *exactly two* horizontal lines (key shapes). The closer the point is to one of the two lines, the more the shape of the (always: "entire") mesh is being influenced by the key shape defined by *that* line, and the correspondingly less it is being influenced by the *other* one.

Options

As with all IPO curves, you can define how the curve is drawn:

- Linear: interpolation between the Keys is linear. The Key line is displayed as a dotted line.
- Cardinal: interpolation between the Keys is fluid, the standard setting.
- BSpline: interpolation between the Keys is extra fluid and includes four Keys in the interpolation calculation. The positions are no longer displayed precisely, however. The Key line is drawn as a dashed line.

Hints

- Key positions are *always* added with I in the 3D View or Materials buttons, even if they are located at the same position. Use this to copy positions when inserting. Two key lines at the same position can also be used to change the effect of the interpolation.
- If *no* Keys are selected, EditMode can be invoked as usual. However, when you leave EditMode, all changes are undone. So if you are frustrated because you edited your mesh and all your changes went away, remember you read this and you will remember why. Insert the Key *in* EditMode in this case to change the active key.
- For Keys, there is *no* difference between *selected* and *active*. It is therefore not possible to select multiple Keys.
- When working with Keys with differing numbers of vertices, the faces can become disordered. There are no tools that can be used to specify *precise* sequence of vertices. This option is actually suitable only for Meshes that have only vertices such as Halos.
- If no IPO curves are defined on any of the absolute Shape Keys, the length of the animation defaults to 100 frames. That is, shape keys added past frame 100 will have no effect. An IPO curve must be

defined on at least the Basis key for absolute shape key animations to exceed 100 frames.

Curve and Surface Keys

As mentioned earlier, Curve and Surface Keys work exactly the same way as Mesh Keys. For Curves, it is particularly interesting to place Curve Keys in the bevel object. Use the animation of that curve to, for example, animate a pumping artery, or a pulsing heart, or a growing worm, or balloon being inflated, or a tree growing. A scene in Elephant's Dream, where worms started growing (just before Proog conks Emo) used animated bevels and tapers while the curves extended.

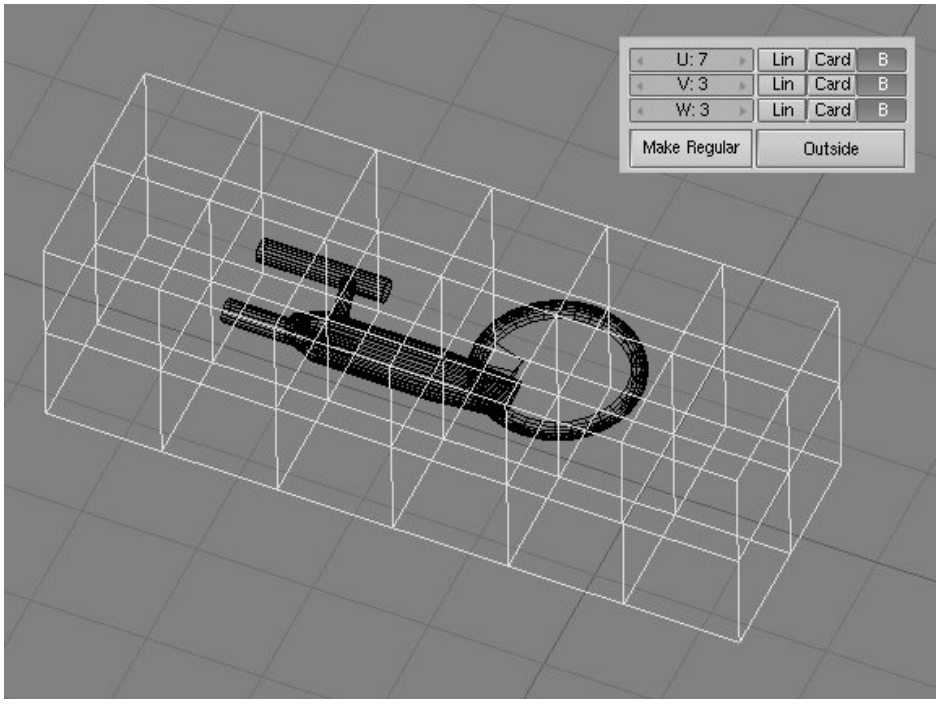
A lattice is a framework of controlling vertices. Associated a mesh to a lattice is a nice way to apply deformations to the mesh while modeling, but it is also a way to make animated deformations in time! To deform the mesh, you can deform the parent Lattice; the mesh deforms to fit inside the Lattice.

Lattice Animation

You can deform a Lattice in animations in a few ways:

- Animate the lattice vertices with shape keys
- Move the child object of the lattice through the lattice over time.

The first technique is basically nothing new than what contained in the previous Shape Key sections but applied to a lattice which has an object parented to it. With the second kind you can create animations that squish things between rollers, or achieve the effect of a well-known space ship accelerating to warp-speed. As the object enters the Lattice, it starts to deform; as it emerges, it expands like a foam rubber ducky back to its original shape.



Lattice setup.
For example, make a space ship mesh object. Add a lattice around the ship. Make the lattice with the parameters shown to the right in *Lattice setup*. This panel is shown in the Buttons window, Editing (F9) context.

There are two ways to make the lattice deform the spaceship:

- Select the object to be deformed by the lattice, and create a Lattice Modifier, entering the name of the Lattice in the Modifier panel.
- Make the object to be deformed a child of the Lattice through a special Parent relationship.

For the old school parenting approach, select the ship, then shift-select (holding â § Shift while selecting) the Lattice, and press CtrlP to make the lattice the parent of the ship. The popup window will ask if you want to make a regular parent relationship, or a Lattice Deform; choose Lattice Deformation.

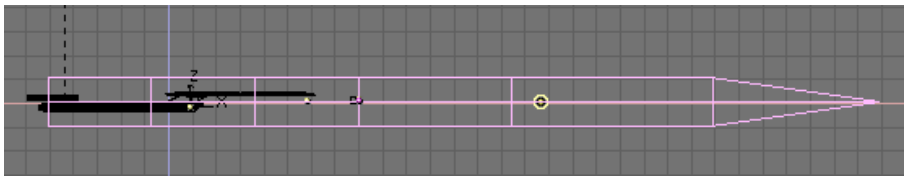
You should not see any deformation of the ship because the lattice is still regular. If you did, click the Regular button in the Lattice panel.

For the next few steps it is important to do them in EditMode. Select the lattice, enter EditMode, and select all vertices (A). Scale the lattice along its x-axis (press MMB while initiating the scale) to get the stretch you want. The ship's mesh shows immediately the deformation caused by the lattice (*Stretching*).



Stretching.
Now edit the lattice in EditMode so that the right vertices have an increasing distance from each other. This will increase the stretch as the ship goes into the lattice. The right ends vertices I have scaled down so that they are nearly at one point; this will cause the vanishing of the ship at the end (*Final lattice deformation*).

Select the ship again and move it through the lattice to get a preview of the animation. Now you can do a normal keyframe animation to let the ship fly through the lattice.



Final lattice deformation.

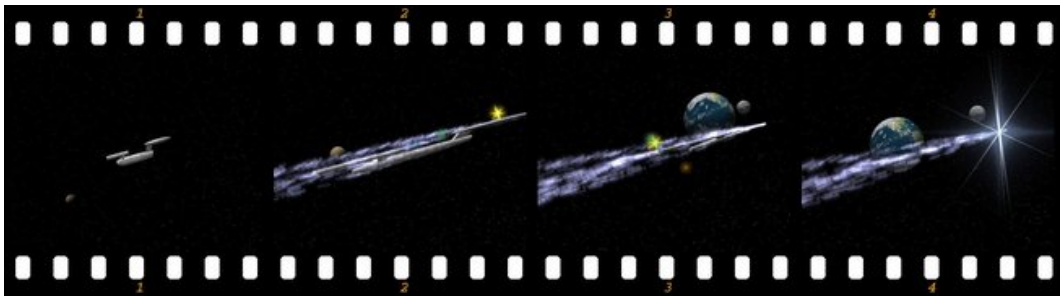
Implications for Camera Tracking

With this lattice animation, you can't use the pivot point of the object for tracking or parenting, since the vertices that make up the mesh are sucked into the Lattice or spit out, way off center from the object itself. When that happens, although the camera is looking at the mesh object's center, the mesh isn't there so the camera is looking in the wrong place. Instead, track an Empty which has an mesh object vertex as a parent.

To do this, create an Empty. Next, make sure the Empty is selected, then select the mesh object. Press â Tab to enter EditMode. Select one vertex from the mesh object. Press CtrlP to make the vertex the parent of the Empty. Press â Tab to get out of EditMode.

Now tell the camera to track to the empty (via the Track To Constraint). Animate the location of the empty to lead the object as it enters the lattice, and lag as it emerges (or vice versa if your Lattice squashes the mesh). As the mesh enters the lattice and deforms, the empty will lead/lag its center, keeping the camera looking at

the deformation action.



Some frames of the resulting animation.
Add some halos and a flash of light (an animated lamp) at the end, and you have something like that shown above.



FIXME

IPO Header



WindowType

As with every window header, the first button enables you to set the window type.

Menus

The triangular button expand/collapses menus. Menus provide a self-explanatory way to access all Blender functions which can be performed in the IPO Window. They are context sensitive and will change depending on the selected Object and current Mode.

Windows with menus, as the IPO Window, does not have the standard "Full Window" and "Home" header buttons, actions which have moved to the View Menu.

Menu items are in general self-explanatory and the relative functionality will be explained later on in the HotKeys section. The actions which do not have an hotkey are the Extrapolation settings :

Extend mode Constant

The ends of selected IPO Curves are horizontally extrapolated.

Extend mode Direction

The ends of selected IPO Curves continue extending in the direction in which they end.

Extend mode Cyclic

The full length of the IPO Curve is repeated cyclically.

Extend mode Cyclic Extrapolation

The full length of the IPO Curve is extrapolated cyclic.

IPO Type



FIXME
Depending on the *active* Object, a **selection** of the following IPO types is available.

Object

Settings, such as the location and rotation, are animated for the *active* Object. All Objects in Blender can have this IPO block. The table below shows the settings that are available for this data block and provides a brief description of each.

IPO Data Elements for Object IPO Data Blocks

IPO Data Item	Description
LocX	Absolute X location
LocY	Absolute Y location
LocZ	Absolute Z location
dLocX	Relative X location
dLocY	Relative Y location
dLocZ	Relative Z location
RotX	Absolute rotation around the X-axis
RotY	Absolute rotation around the Y-axis
RotZ	Absolute rotation around the Z-axis
dRotX	Relative rotation around the X-axis
dRotY	Relative rotation around the Y-axis

dRotZ	Relative rotation around the Z-axis
ScaleX	Absolute scaling along the X-axis
ScaleY	Absolute scaling along the Y-axis
ScaleZ	Absolute scaling along the Z-axis
dScaleX	Relative scaling along the X-axis
dScaleY	Relative scaling along the Y-axis
dScaleZ	Relative scaling along the Z-axis
Layer	Changes in Layer
Time	Changes in time
ColR	Changes to the Red component of the color
ColG	Changes to the Green component of the color
ColB	Changes to the Blue component of the color
FStrength	Strength of Force Fields
FFall	Fall-off of Force Fields
RDamp	Collision Damping
Damping	Random Variation of Collision Damping
Perm	Deflector Permeability - percentage of particles allowed to pass deflector mesh

Material

Settings of the *active* Material are animated for the *active* Object. A numeric field appears immediately to the right of the IPO Type selection field when Material is selected. This field indicates the number of the active Texture *channel*. Eight Textures, each with its own mapping, can be assigned per Material. Thus, per Material-IPO , 8 curves in the row OfsX, OfsY,...Var are available. To change which material you are working on, change the number in the field by left clicking on the field and typing in a new value from 0 to 7.

IPO Data Elements for Material IPO Data Blocks

IPO Data Item	Description
R	Red
G	Green
B	Blue
SpecR	Red component of the specular reflection
SpecG	Green component of the specular reflection
SpecB	Blue component of the specular reflection
MirR	Red component of mirrored light
MirG	Green component of the mirrored light
MirB	Blue component of the mirrored light

Ref	The amount of reflection for diffuse shader
Alpha	Alpha value for the texture
Emit	Sets the amount of light the material emits
Amb	Amount of global ambient color the material receives
Spec	Degree of specularity
Hard	Hardness setting for specular shader
SpTra	SpecTra - makes specular areas opaque on translucent materials
Ior	Index of Refraction
Mode	These are quite a lot of different material settings (I've counted 22 different settings, from shadeless to halo to colorband and so forth).
HaSize	Halo Size
Translu	Translucence - diffuse shading of the back side
RayMir	Amount of mirror reflection for ray trace
FresMir	Power of Fresnel for mirror reflection
FresMirl	*Fac*-Value for Fresnel mirror
FresTra	Fresnel Transparency
FresTral	*Fac*-Value for Fresnel transparency
TraGlow	*Add*-Value for Halo material
OfsX	Fine tunes the X mapping coordinates (Map Input)
OfsY	Fine tunes the Y mapping coordinates (Map Input)
OfsZ	Fine tunes the Z mapping coordinates (Map Input)
SizeX	Sets scaling on the texture's X-axis (Map Input)
SizeY	Sets scaling on the texture's Y-axis (Map Input)
SizeZ	Sets scaling on the texture's Z-axis (Map Input)
texR	Sets Red component of the Map To color
texG	Sets Green component of the Map To color
texB	Sets Blue component of the Map To color
DefVar	Default value the texture uses to mix with other textures
Col	Sets the amount the texture affects color values
Nor	Sets the amount the texture affects the normal
Var	Sets the amount the texture affects other values
Disp	Sets the amount the texture displaces the surface

World

Used to animate a number of settings for the WorldButtons. World too has several texture channels.

IPO Data Elements for World IPO Data Blocks

IPO Data Item	Description
HorR	Red component of the horizon color (world)
HorG	Green component of horizon color (world)
HorB	Blue component of horizon color
ZenR	Red component of the color at the zenith
ZenG	Green component of the color at the zenith
ZenB	Blue component of color at the zenith
Expos	Exponential color correction for the light
Misi	Controls Mist Simulation (mist/stars/physics)
MisDi	Mist Distance (mist/stars/physics)
MisSta	Start of Mist (mist/stars/physics)
MisHi	Height of Mist (mist/stars/physics)
StarR	Amount of red in star color
StarG	Amount of green in star color
StarB	Amount of blue in star color
StarDi	Star distance
StarSi	Star size
OfsX	Offset of texture on x-axis
OfsY	Offset of texture on y-axis
OfsZ	Offset of texture on z-axis
SizeX	Sets scaling on the texture's X-axis (Map Input)
SizeY	Sets scaling on the texture's Y-axis (Map Input)
SizeZ	Sets scaling on the texture's Z-axis (Map Input)
texR	Sets Red component of the Map To color
texG	Sets Green component of the Map To color
texB	Sets Blue component of the Map To color
DefVar	Default value the texture uses to mix with other textures
Col	Sets the amount the texture affects color values
Nor	Sets the amount the texture affects the normal
Var	Sets the amount the texture affects other values

Shape Key

NOTE: Shape keys used to be known as Vertex Keys and Relative Vertex Keys, or RVKs.

Each key applied to the object will be displayed as a field in the right hand portion of the IPO editing window, like the fields for other IPO types. The curves created here are used to control the amount of influence each of the keys has over the shape.

See Absolute Shape Keys and Shape Keys for information on how shape keys are created and how they can be modified using the IPO Data blocks.

Constraint

If the *active* Object has a *constraint* its influence value can be animated via an IPO. Each constraint has its own IPO. Used to display the speed-IPO.

Sequence

The *active* Sequence Effect can have an IPO Curve.

Curve IPO

If the *active* Object is a *path* Curve, this button can be used to display the speed (time) IPO.

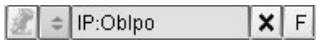
Camera IPO

The active camera IPO curves are shown.

Lamp IPO

If the *active* Object is a Lamp, this button can be used to animate light settings, comprehensive of textures.

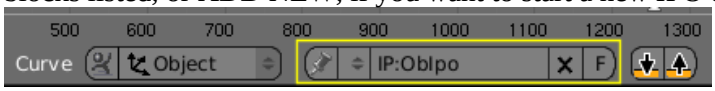
IPO Menu



FIXME
The DataButtons can be used to control which IPO block is shown and control it.



IPO Window Bar when there is no preexisting IPO data block.
NOTE: When there is no IPO data block associated with an object, the toolbar will look like this. Instead of an IPO data block reference, like you see in the second image, there is only the up/down arrow button, as highlighted in the graphic. To add a data block, just click on the button and select one of the existing data blocks listed, or ADD NEW, if you want to start a new IPO data block.



IPO Window Bar With Existing Data Block
When there is an existing IPO block for an object, the toolbar for the IPO Window will look like the highlighted section in the second graphic.

Pin IPO

The IPO Window shows the current IPO even if the linked Object is deselected.

IPO Menu

Choose another IPO from the list of available IPOs. The option Add New makes a complete copy of the current IPO. This is not visible; only the name in the adjacent button will change. Only IPOs of the same type are displayed in the menu list.

IP :

Give the current IPO a new and unique name. After the new name is entered, it appears in the list, sorted alphabetically.

Users

If this button is displayed, there is more than one user for the IPO block. Use the button to make the IPO "Single User ".

Unlink IPO

The current IPO is unlinked.

Fake User



FIXME
The IPO block is saved even if unused.

Copy to Buffer

All selected IPO Curves are copied to a temporary buffer.

Paste from Buffer

All selected *channels* in the IPO Window are assigned an IPO Curve from the temporary buffer. The rule is: the sequence in which they are copied to the buffer is the sequence in which they are pasted. A check is made to see if the number of IPO Curves is the same.

```
NOTE: If you want to paste IPO data onto an object, the object must
   have an IPO data block to begin with.
```

View Border

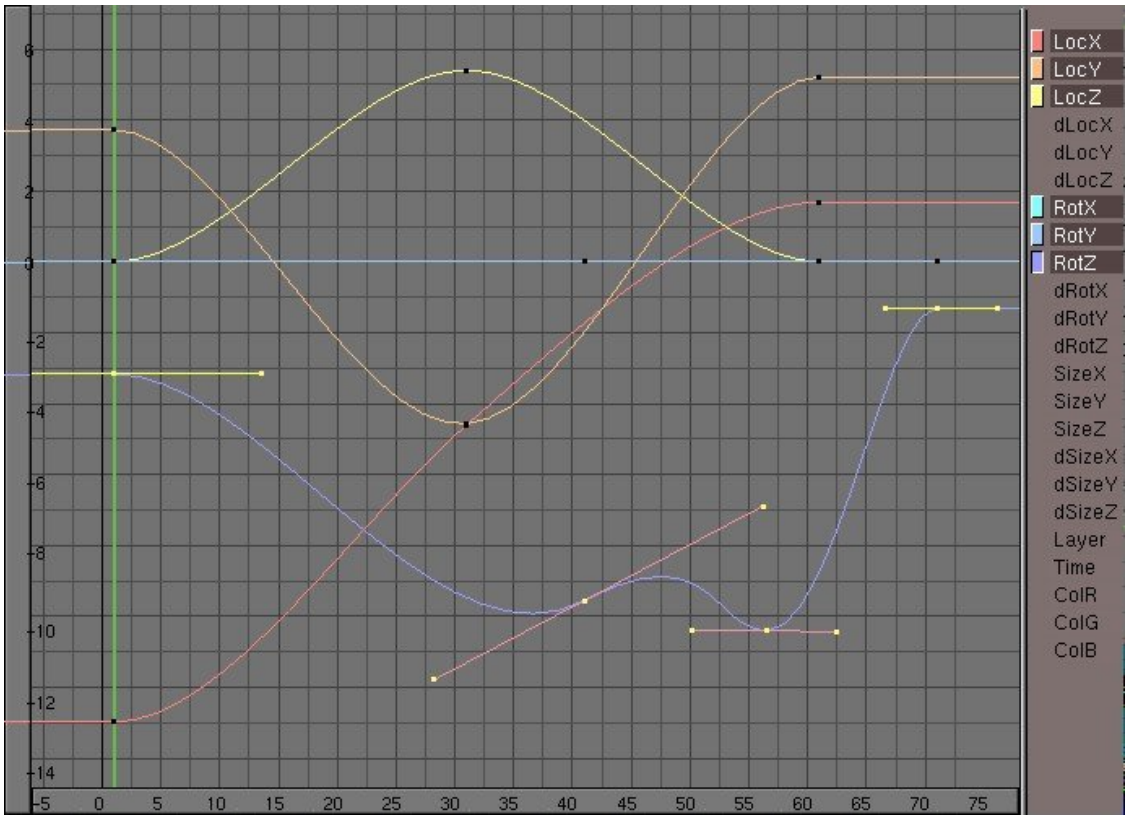


FIXME
Draw a rectangle to indicate what part of the IPO Window should be displayed in the full window.

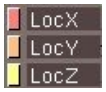
Lock

This button locks the update of the 3DWindow while editing in the IPO Window, so you can see changes made to the IPO in realtime in the 3DWindow. This option works extremely well with relative vertex keys.

IPO Window



FIXME



FIXME
The IPO Window shows the contents of the IPO block. Which one depends on the IPO Type specified in the header. The standard IPO Window has a grid with the time expressed horizontally in frames and vertical values that depend on the *channel*. There are 2 sliders at the edge of the IPO Window. How far the IPO Window is zoomed in can be seen on the sliders, which can also be used to move the *view* .

The right-hand part of the window shows the available *channels* . To make it easier to work with rotation-IPO Curves, they are displayed in degrees (instead of in radiants). The vertical scale relation is: `1.0 'Blender unit' = 10 degrees`. In addition to the IPO Curves, the VertexKeys are also drawn here. These are horizontal blue lines; the yellow line visualises the *reference* Key.

Each *channel* can be operated with two buttons :

IPO Curve Select

This button is only displayed if the *channel* has an IPO Curve. The button is the same colour as the IPO Curve. Use the button to select IPO Curves. Multiple buttons can be (de)selected using â § ShiftLMB .

Channel Select

A *channel* can be selected whether there is an IPO Curve or not. Only IPO Curves of selected channels are drawn. Multiple *channels* can be (de)selected using â § ShiftLMB .

Mouse

CtrlLMB

Create a new vertex. These are the rules :

- There is no IPO block (in this window) *and* one *channel* is selected: a new IPO Block is created along with the first IPO Curve with one vertex.
- There is already an IPO block, and a *channel* is selected without an IPO Curve: a new IPO Curve with one vertex is added.
- Otherwise a new vertex is simply added to the selected IPO Curve.
- This is *not* possible if multiple IPO Curves are selected or if you are in EditMode.

MMB

Depending on the position within the window :

- On the *channels*; if the window is not tall enough to display them completely, the visible part can be scrolled up and down.
- On the sliders; these can be moved. This only works if you are zoomed in.
- The rest of the window; the *view* is translated.

CtrlMMB

Zoom in/out on the IPO Window. You can zoom horizonally or vertically using horizontal and vertical mouse movements.

RMB

Selection works the same here as in the 3DWindow: normally one item is elected. Use â § Shift to add/remove from the selection.

- If the IPO Window is in IPO Key mode, the IPO Keys can be selected.

- If at least 1 of the IPO Curves is in EditMode, only its vertices can be selected.
- VertexKeys can be selected if they are drawn (horizontal lines).
- The IPO Curves can be selected.

RMB and drag

Select and start *translation* mode, i.e. the Grabber. The selection can be made using any of the four selection methods discussed above.

â § ShiftRMB

Adds/removes from the selection.

HotKeys

- NumPad + NumPad

Zoom in, zoom out.

Pageup

Select the next IPO Key. If more than one IPO Key is selected, the selection is shifted cyclically.

â § ShiftPageup

Add the next IPO Key to the selection.

Pagedown

Select the previous IPO Key. If more than one Object Key is selected, the selection is shifted cyclically.

â § ShiftPagedown

Add the previous IPO Key to the selection.

â Home

All visible curves are displayed completely, centered in the window.

â Tab

All selected IPO Curves go into or out of EditMode. This mode allows you to transform individual vertices.

A

Select All / deselect All. If any item is selected, first everything is deselected. Placing the mouse cursor above the *channels*, (de)selects all channels where there is a curve.

B

Border select. Draw a rectangle with the LMB; all items that fall within this rectangle are *selected*. Draw a rectangle with the RMB to *deselect* .

â § ShiftC

Centers the view on the current *frame*, without zooming.

C

If one vertex or one IPO Key is selected, the current *frame* number is set to this position.

â § ShiftD

Duplicate IPO. All selected vertices or IPO Keys are copied. Then *translation* mode is started automatically.

G

Translation mode (the Grabber). This works on selected curves, keys or vertices. Alternatives for starting this mode :

- RMB and drag.

The following options are available in *translation* mode:

- Limiters:
- Ctrl increments of 1 frame or vertical unit.
- â § ShiftCtrl increments of 0.1 frame or vertical unit.
- MMB restricts the current translation to the X or Y axis. Blender calculates which axis to use, based on the already initiated mouse movement. Click MMB again to restore unlimited translation.
- ←↑↓→: With these keys the mouse cursor can be moved exactly 1 pixel.
- X Limits translation to the X axis.
- Y Limits translation to the Y axis.

Grabber can be terminated with:

- LMB, Space or â µ Enter: Move to the new position.

- RMB or Esc: Everything returns to the old position.

H

Toggle Handle *align* / *free* .

â § ShiftH

Set Handle *auto*. The selected Bezier *handles* are converted to *auto* type.

I

Insert Key. Vertices can be added to the visible curves in the IPO Window. A PopupMenu asks you to make a choice :

- Current Frame: All visible curves get a vertex on the current frame.
- Selected Keys: (only in IPO Key mode) all selected IPO Keys get vertices for each visible curve, including IPO Curves that are not part of the IPO Key.

J

Join vertices. Selected vertices or IPO Keys can be joined. A PopupMenu asks you to make a choice :

- All Selected: all selected vertices are replaced by a new one.
- Selected doubles: all selected vertices that are closer to each other than 0.9 *frame* are joined.

K

IPO Key mode ON/OFF. If the IPO block is Object IPO type, the Objects are redrawn with the option DrawKey ON (see the explanation under IPO Header).

R

Recording mode. The X and Y movements of the mouse are linked to the height of the IPO Curve. Thus, this works with a maximum of two selected *channels* or IPO Curves. The old curve is completely deleted; the new one becomes a 'linear' type. You cannot change parts of curves with *recording*. The scale at which this happens is determined by the *view* of the IPO Window. A PopupMenu asks you to make a choice :

- Still: The current frame is used as the starting point.
- Play anim: The animation starts, allowing you to see the correlation with other animation systems.

During *recording* mode, the Ctrl must be held down to actually start recording. Press Space or â µ Enter or LMB to stop *recording*.

Use Esc to undo changes.

S

Scaling mode - This works on selected IPO Curves and vertices. The degree of *scaling* is *precisely* linked to the mouse movement. Try to move from the (rotation) midpoint with the mouse.

In IPO Key mode, you can only *scale* horizontally.

Limiters :

- Ctrl: in increments of 0.1.
- â § ShiftCtrl: in increments of 0.01.
- MMB limits *scaling* to the X or Y axis. Blender calculates which axis to use based on the already initiated mouse movement. Click MMB again to return to free *scaling*.
- ←↑↓→: These keys allow you to move the mouse cursor exactly 1 pixel.
- X Limits *scaling* to the X axis.
- Y Limits *scaling* to the Y axis.

Terminate size mode with:

- LMB Space or â µ Enter: To finalize scaling.
- RMB or Esc: Everything returns to its previous state.

â § ShiftS

Snap Menu.

- Horizontal: The selected Bezier handles are set to horizontal.
- To next: The selected handle or vertex is set to the same (Y) value as the next one.
- To frame: The selected handles or vertices are set to the exact frame values.
- To current frame: The selected handle or vertex is moved to the current frame.

T

If an IPO Curve is selected: "IPO Type". The type of selected IPO Curves can be changed. A PopupMenu asks you to make a choice :

- Constant: After each vertex of the curve, this value remains constant, and is not interpolated.
- Linear: Linear interpolation occurs between the vertices.
- Bezier: The vertices get a *handle* (i.e. two extra vertices) with which you can indicate the curvature of the interpolation curve.

If a Key is selected: "Key Type". The type of selected Keys can be changed.

- Linear: Linear interpolation occurs between the Keys. The Key line is displayed as a broken line.
- Cardinal: Fluent interpolation occurs between the Keys; this is the default.
- BSpline: Extra fluent interpolation occurs between the Keys, four Keys are now involved in the interpolation calculation. Now the positions *themselves* cannot be displayed precisely, however. The Key line is shown as a broken line.

V

Vector Handle. The selected Bezier *handles* are converted to *vector* type.

X

Erase selected. The selected vertices, IPO Keys or IPO Curves are deleted. If there are selected VertexKeys, they are also deleted.

See also

- PART VII - ANIMATION BASICS

Hooks

Hooks are similar to shape keys in that they deform a mesh over time (frames). The difference is that hooks make it look like the mesh is snagged with a fish hook. Moving the hook moves selected vertices under the influence of the hook (which is usually just an empty â← but might be any object), and you make the hook move by animating the motion of the empty through Ipo keys or any other animation method. As the hook moves, it weighted-pulls vertices from the mesh with it. If you have used proportional editing, you can think of it as animated proportional editing. While hooks do not give you the fine control over vertices movement that shape keys do, they are much simpler to use.

Object Hooks

Mode: Object mode or Edit mode

Panel: Modifiers (Editing context, F9)

Hotkey: CtrlH

Menu: Mesh Â» Vertices Â» Add Hook

Description

Hooks give access at object level to the underlying geometry of meshes, curves, surfaces or lattices. A hook is an object feature and it is like a parent to an object, but for vertices. You can create as many hooks to an object as you like, and assign for each hook vertices that will be affected. Overlapping hooks is also possible, here a weighting factor per hook is provided that determines the amount each hook will affect the overlapping vertices.

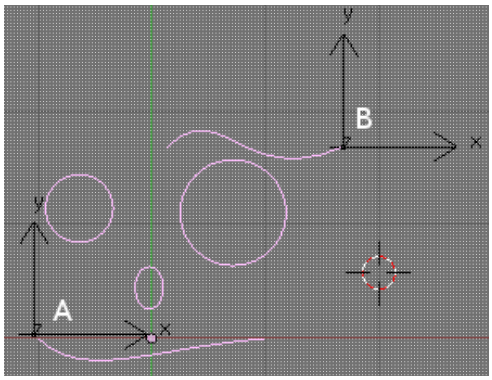
All object level options and transformations are possible on the hook object, including using hierarchies, constraints, Ipo and path animations. You can also make the hook-parent a child of the original object if you donâ← t want object transformations to deform the hooks.

Note
When you change topology (i.e. more destructive editing than just manipulating existing points), you most likely have to reassign existing hooks as well.

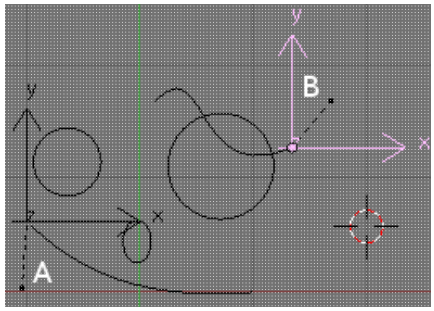
Examples

A typical example of using hooks is to animate vertices or groups of vertices. For example, you may want to animate the vertices associated with a â← Mouthâ← on a characterâ← s face.

In (*Animated face frame 1*) and (*Animated face frame 10*) a face made from BÃ©zier curves has two hooks applied. One applied to a control-point on the mouth labelled â← Aâ← and one applied to the eyebrow labelled â← Bâ← . The animation has 10 frames over which hook A moves up and hook B moves down.



Animated face frame 1.



Animated face frame 10.

Adding Hooks

Mode: Edit mode

Panel: Modifiers (Editing context, F9)

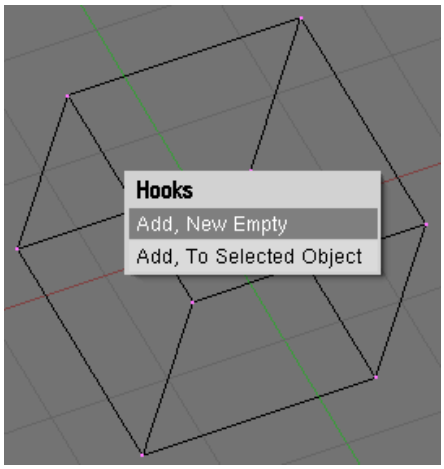
Hotkey: CtrlH

Menu: Mesh Â» Vertices Â» Add Hook

Description

Since hooks relate to vertices or control points, most editing options are available in Edit mode for meshes, curves, surfaces and lattices (not texts, as you might have guessed!).

Options



Hooks menu.

Add, New Empty
: Adds a new hook and creates a new empty object, which will be the hook controller, at the center of the selection.

Add, To Selected Object
: When another object is selected (you can do this in Edit mode with CtrlRMB), the new Hook modifier is created with that selected object as hook controller.

Editing Hooks

Mode: Edit mode

Panel: Modifiers (Editing context, F9)

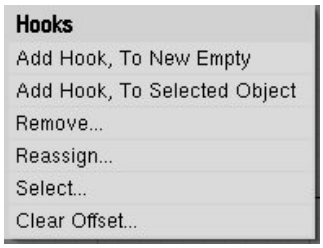
Hotkey: CtrlH

Menu: Mesh Â» Vertices Â» Add Hook

Description

Once hooks are available in an object, the Hooks menu will give additional options:

Options



Extended Hooks menu.

Remove...
: This will give a new menu with a list of hooks to remove.

Reassign...
: Use this if you want to assign new vertices to a hook.

Select...
: Select the vertices attached to a specific hook

Clear Offset...
: Neutralize the current transformation of a hook parent.

Hook Modifier

Mode: Object mode or Edit mode

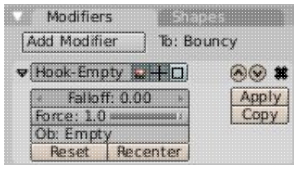
Panel: Modifiers (Editing context, F9)

Hotkey: CtrlH

Description

Hooks are modifiers, that are added to the modifier stack. For each Hook modifier, you can give a hook a new name (the default name is the hook controller name, prefixed with â← `Hook-`â←), give it a new controller, or assign it a Force weighting factor.

Options



In the Editing context (F9), Modifiers panel, when a hook is created, you can control it via this panel.

Ob
: The parent object name for the hook. Changing this name also recalculates and clears offset.

Reset
: Recalculate and clear the offset transform of hook.
The same as the Clear Offset... option from Hooks menu (CtrlH).

Recenter
: Set hook center to cursor position.

Select â← Edit mode only
: Select affected vertices on mesh.
The same as the Select... option from Hooks menu (CtrlH).

Reassign â← Edit mode only
: Reassigns selected vertices to this hook.
The same as the Reassign... option from Hooks menu (CtrlH).

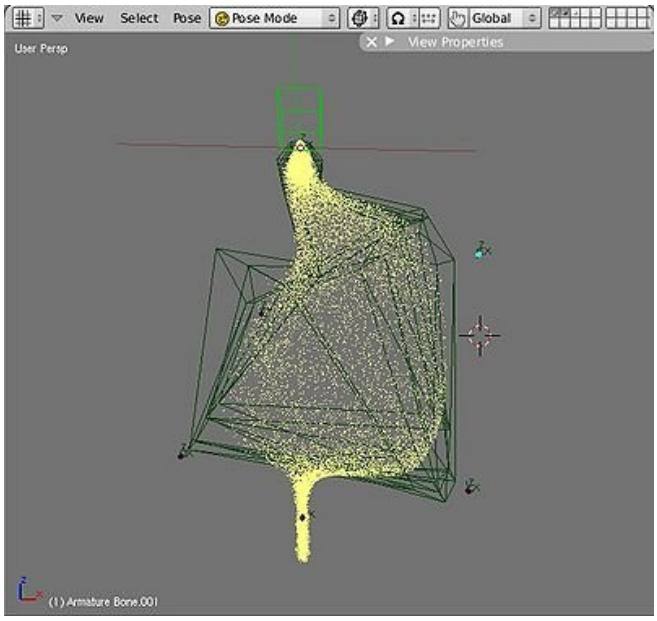
Force
: Since multiple hooks can work on the same vertices, you can weight the influence of a hook this way.
Weighting rules are:
 - ◊ If the total of all forces is smaller than **1.0**, the remainder (`1.0 - (forces sum)`, will be the factor the original position have as force.
 - ◊ If the total of all forces is larger than **1.0**, it only uses the hook transformations, averaged by their weights.

Falloff

If not zero, the falloff is the distance where the influence of a hook goes to zero. It currently uses a smooth interpolation, comparable to the proportional editing tool.

Examples

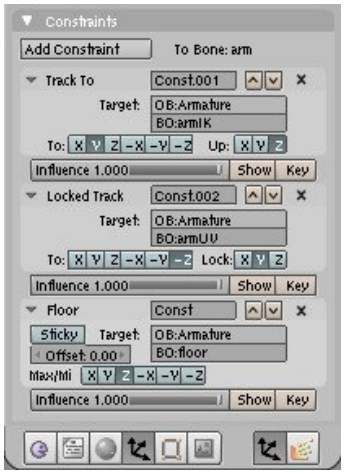
Hook modifier on lattices



Introduction To Blender Constraints

Constraints are object features that define spatial relationships between objects, and are the standard method for controlling characters among all 3D animation packages that still implement a more traditional approach to digital character animation. In Blender, constraints can be associated to any type of object or bone object, but not all constraints work with bone objects, and not all constraints work with normal world objects.

General Constraint Properties



Constraints are accessed via the object buttons (F7) in the Constraints panel. After you press the Add Constraint button and select the desired constraint type from the menu, a constraint UI is added to the panel. The constraint will be linked to the active object or the active selected bone, indicated by the label To Object: or To Bone: on the right of the Add Constraint menu.

Constraints are evaluated in a specific order: from top to bottom of the constraint stack. This order can be viewed and changed inside the Constraints panel. Each constraint has a pair of arrow buttons on the right of the constraint name in its top right corner, which can be used to move the constraint up or down the constraint stack. The little cross on the right of these buttons can be used to delete the constraint from the stack.

NOTE: The name field of a newly added constraint will appear red, which indicates that the constraint is not functional. In the case of a newly added constraint this is because the constraint does not yet have a target specified. All constraints require a target object (a normal world object or an armature bone). You should enter the name of the desired target object in the Target: OB field. When you want an armature bone as target, enter the name of the armature object as target. A new field with BO: will appear, where you can place the name of the target bone.

The name field of a constraint will also turn red if the settings for the constraint are invalid or if the constraint conflicts with another constraint. For example: A Track To constraint of which the To and Up vector are set to Z.

Influence

The **influence** of a constraint on the actual location/rotation/size of the object can be defined with the Influence value slider. This can be linked to an Ipo Curve, which can be evoked in the Ipo Curve Editor with the button Show on the right of the Influence value slider. The Key button beside it can be used to add a key to the Ipo Curve.

You can key any constraint, and an object can have multiple constraints keyed to exert a different influence at different times.

Multiple Influence Example

Suppose that you want a camera to follow one path for awhile and then drift off and follow another path. You apply a *follow path* constraint to a camera for one circle/curve/path (make sure you have curve path enabled and enter its name), then open up an IPO curve window. In the buttons window (constraint panel) you can either hit I and insert an Influence value by sliding the slider (to vary the force strength or force fall off) or LMB click the "key" button at the bottom of the constraint panel. Once you have inserted the initial key, up arrow or change to the frame where you want the first influence to start falling off. Go to the IPO window and and CtrlLMB to insert a key. Go forward a few frames, where you want its influence to be zero, and insert another key, and edit it so that its Y value is zero. At that frame, the constraint in the stack will not influence the camera motion at all. Now insert a *second* Follow Path constraint, and enter the name of the *second* path. Now position back a few frames and insert a key to where you want its influence to start (probably where the first constraint starts dropping off). Edit this second influence curve to come up from zero to one over time. You can imagine that the IPO curves for each circle should cross over each other so that when one is at full influence (max 1.0) the other is at zero etc. You might have to play with the curves a bit to get the transition smooth.

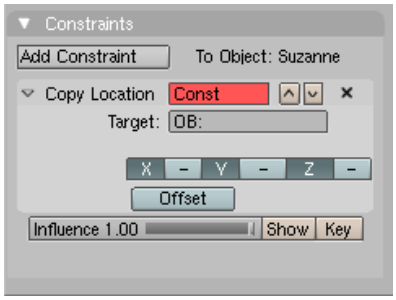
Constraints on Bones

Concerning bones with constraints: the color of the bones in the 3D Window indicate what type of constraint is used:

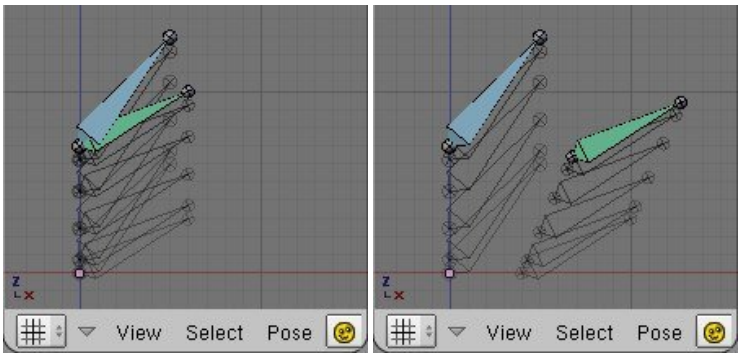
- **Grey**: No constraint.
- **Yellow**: A bone with an IK constraint.
- **Orange**: A bone with an IK constraint but no target.
- **Green**: A bone with any other kind of constraint.
- **Blue**: A bone that is animated with keyframes.
- **Purple**: The Stride Root.

Constraint Types

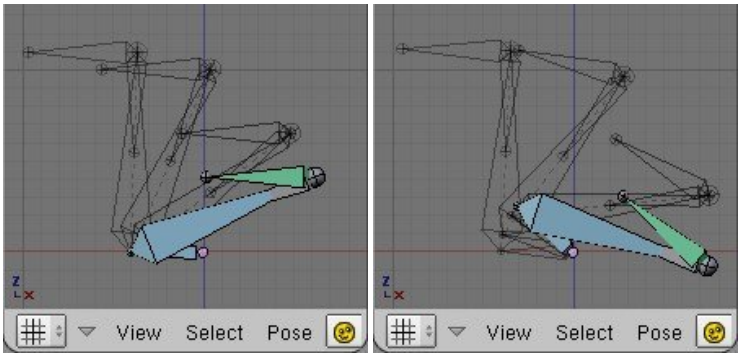
Copy Location



Copy Location forces the object to have the same location as its target. When this constraint is used on a bone and another bone is the target, a new button--labeled **Local**--appears next to the **X Y Z** buttons. Using the local button causes the local translation values of the target to be given to the constrained object. In rest position, all bones are at the local location of (0,0,0).

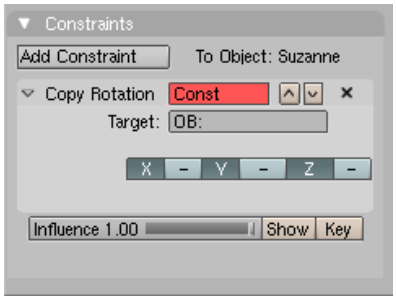


Left: Using global space. This is the default behavior; it's what happens when the local button is not activated. Right: Using local space, with local button activated.

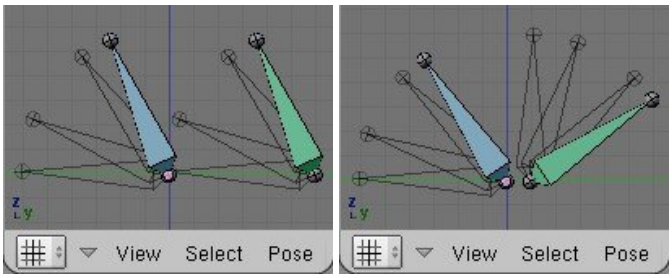


In these last two images, the constrained bone (shown in green) is actually a child of the root bone (the bone at the beginning of the chain). This demo shows possible uses for the location constraint in a rig. Note that the green bone still inherits rotation from the root because the root is its parent. This is by design though; the green bone could be the child of any of these bones, or none of them.

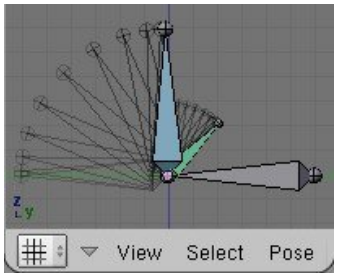
Copy Rotation



Copy rotation causes one object to match the rotation of a target object. For bones, a local option will appear, allowing you to use local space. Imagine you have a bone pointing up and a bone pointing down. If you use local space, each can point different directions, but when the target bone moves to its left, the affected bone will move it to *its* left.

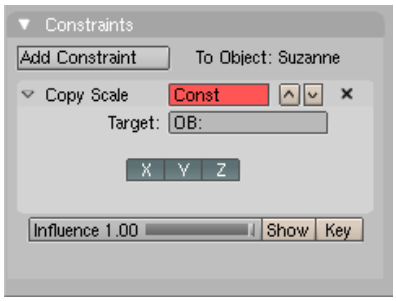


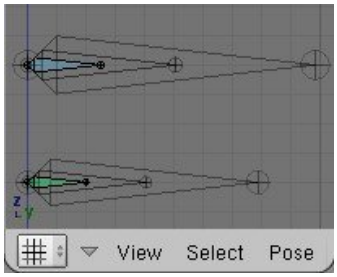
Left: Using global space. Right: Using local space.



Here is one good use of the rotation constraint and local space. By setting the influence value to 0.5, the small bone will rotate half as far as the target. This is useful for character joints.

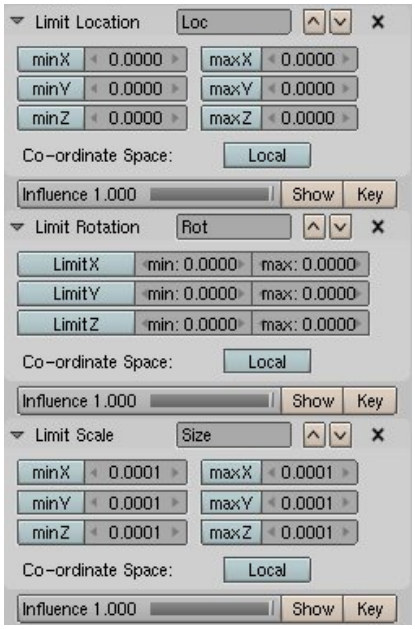
Copy Scale





Copy Scale forces the affected object to have the same size as the target. All bones have a (1,1,1) size in rest position. We can draw a bone that is 10 million times longer than the bone right next to it, but in pose mode, they both have a starting size of (1,1,1). You should keep this in mind if you're going to be using the copy scale constraint.

Limit Constraints



Limit Loc/Rot/Size constraints
These constraints limit the scale, rotation, or location of the active object. The limit can be named, and is always specified in XYZ terms. Each limit min or max must be enabled by LMB, turning it dark green.

Limit Location
: The object cannot be located or positioned below the minimum, or above the maximum amount. Use this to constrain an object to only move within a confined space.

Limit Rotation
: The object cannot be rotated below the minimum, or above the maximum amount. Use this to constrain a bone, for example, not to hyperextend.

Limit Scale
: The object cannot be scaled below the minimum, or above the maximum amount.

Track To

The Track To constraint applies rotations to the object, so that it always points a local â← Toâ← axis towards the target object, with another local â← Upâ← axis permanently maintained as vertical as possible. This tracking is similar to the billboard tracking in 3D. This is the preferred tracking constraint, because it has a more easily controlled constraining mechanism.

The â← billboard trackingâ←
The term â← billboardâ← has a specific meaning in real-time CG programming (i.e. video games!), where it is used for plane objects always facing the camera (they are indeed â← trackersâ← , the camera being their â← targetâ←). Their main usage is as support for tree or mist textures: if they werenâ← t permanently facing the camera, you would often see your trees squeezing to nothing, or your mist turning into a millefeuille paste, which would be funny but not so credibleâ←¦

This constraint shares a close relationship to the IK constraint in some ways. This constraint is very important in rig design, and you should be sure to read and understand the tracking tutorial, as it centers around the use of both of these constraints.

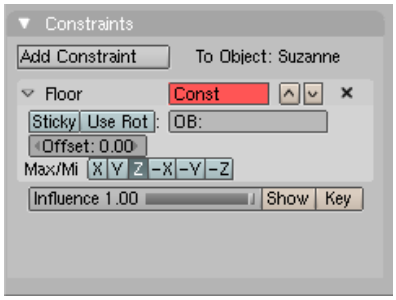
Settings

To
: The axis of the object that has to point to the target.

Up
: The axis of the object that has to be aligned (as much as possible) with the world Z axis. An Align: Target button, when enabled, uses the coordinates of the target object's Z axis, thus tilting or rocking the object as it tracks the target.

Head/Tail - With Bone targets only
: A number from 0.0 to 1.0 that represents the place on the target bone to Track to (0.0 = the bone's root; 1.0 = the bone's head)

Floor



Animation Tip:

When you animate foot placement on the floor plane, ***always*** be sure to use the option **VisualLoc** from the **Insert Key** menu, or enable **Use Visual Keying** from the **Auto Keyframing** menu in the User Preferences.

The Floor Constraint allows you to use a target object to specify the location of a plane which the affected object cannot pass through. In other words, it creates a floor! (or a ceiling, or a wall). This only works by default with planes of the global coordinate system.

Settings

Sticky
: Makes the affected object immovable when touching the plane (cannot slide around on the surface of the plane), which is fantastic for making walk and run animations.

UseRot
: Takes the target object's rotation into account.

Offset
: A number of BU's from the object's center. Use this to account for the distance from a foot bone to the surface of the foot's mesh.

Max/Min
: Which axis is the floor? Things normally stick "down" Z, but can walk on walls as well.

Head/Tail - With Bone targets only
: A number from 0.0 to 1.0 that represents the place on the target bone to use for the Floor (0.0 = the bone's root; 1.0 = the bone's head)

Locked Track

Locked Track is a difficult constraint to explain, both graphically and verbally. The best real-world example would have to be a compass. A compass can rotate to point in the general direction of its target, but it canâ← t point *directly at* the target, because it spins like a wheel on an axle. If a compass is sitting on a table and there is a magnet directly above it, the compass canâ← t point to it. If we move the magnet more to one side of the compass, it still canâ← t point *at* the target, but it can point in the general direction of the target, and still obey its restrictions of the axle.

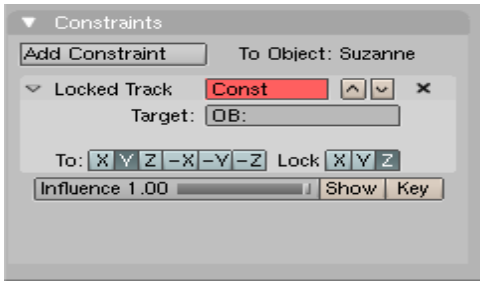
When using a Locked Track constraint, you can think of the target object as a magnet, and the affected object as a compass. The â← Lockâ← axis will function as the axle about which the object spins, and the â← Toâ← axis will function as the compass needle. Which axis does what is up to you! If you have trouble understanding the buttons of this constraint, read the tool-tips; theyâ← re pretty good. If you donâ← t know where your objectâ← s axes are, turn on the Axis button in the Draw panel, Object context (F7). Or, if youâ← re working with bones, turn on the Draw Axes button, Armature panel, Editing context (F9).

This constraint was designed to work cooperatively with the Track To constraint. If you set the axes buttons right for these two constraints, Track To can be used to point the axle at a target object, and Locked Track can spin the object around that axle to a secondary target.

This constraints also works very well for 2D billboarding.

This is all related to the topic discussed at length in the tracking tutorial.

Settings



The Locked Track Constraints panel.

Target
: The name of the target object that the tracking object tracks.
VG
: This is only available when tracking to a mesh object. If you type in the name of a vertex group of the target, it will be the center of this vertex group that will be the â← physicalâ← target, not the objectâ← s center.
BO
: This is only available when tracking to an armature object. If you type in the name of a bone of the target, it will be this bone that will be the â← physicalâ← target, not the armatureâ← s center.

To
: The tracking axis. It shouldnâ← t be the same as the Lock axis.

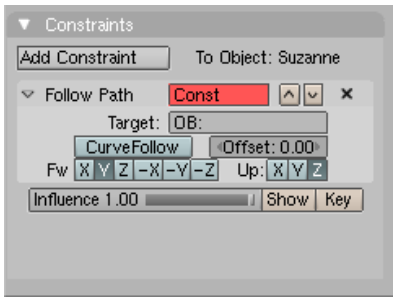
Lock
: The locked local axis.

Influence
: This controls how accurately the tracking object tracks the target. **0** means that the constraint is turned off. The tracking object will remain locked in orientation. **1** means tracking is completely on and the tracking axis will stay tightly focused on the target.

Show
: This adds an influence Ipo channel to the constraint if one is not present. You can then add keys to the channel.

Key
: This adds animation keys to the influence Ipo channel. This is a very powerful combination. For example, you could have a camera with a Locked Track constraint applied and have input driving the influence channel.

Follow Path



Follow Path places the affected object onto a curve object. Curves have an animated property that causes objects along the path to move. In order for this to work, you have to have the CurvePath option activated (it's a property of the curve object). You can find it in the Curve and Surface panel in the Editing buttons, F9.

The movement along the path might be controled by two different ways: the most simple, in this same Curve and Surface panel, is to define the number of frames of the movement *via* the num button Path Len:, and its start frame *via* the constraint's Offset: option (by default: start frame **1** (= offset of **0**), duration **100**).

The second way much more precise and powerful is to define a Speed IPO curve for the path (Path section of the IPO Curve window). The start position along the path will correspond to an IPO value of **0.0**,

and the end position, to an IPO value of **1.0**. You can therefore control the start frame, the speed of the movement, and the end frame, and even force your object to go forth and back along the path!

If you don't want objects on the path to move, you can give the path a flat speed IPO curve (its value will control the position of the object along the path).

Follow Path is another constraint that works well with Locked Track. One example is a flying camera on a path. To control the camera's roll angle, you can use a Locked Track and a target object to specify the up direction, as the camera flies along the path.

This constraint does not work well with bones.

Settings

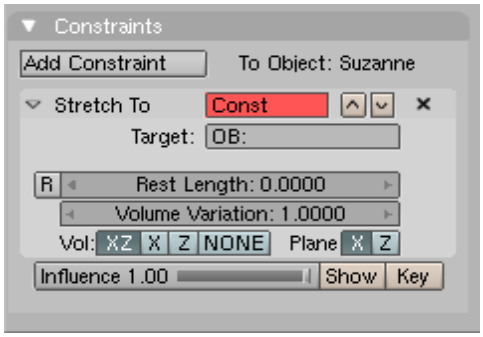
Offset
: The number of frames to offset from the "animation" defined by the path (by default: from the frame **1**).

CurveFollow
: If this option isn't activated, the affected object's rotation isn't modified by the curve; otherwise, it's affected depending on the following options:

Fw
: The axis of the object that has to be aligned with the forward direction of the path.

Up
: The axis of the object that has to be aligned (as much as possible) with the world Z axis.
In fact, with this option activated, the behaviour of the affected object shares some properties with the one caused by a Locked Track constraint, with the path as "axle", and the world Z axis as "magnet"

Stretch To



Stretch To causes the affected object to scale the Y axis towards a target object. It also has volumetric features, so the affected object can squash down as the target moves closer, or thin out as the target moves farther away. Or you can choose not to make use of this volumetric squash-'n'-stretch feature, by pressing the **NONE** button. This constraint assumes that the Y axis will be the axis that does the stretching, and doesn't give you the option of using a different one because it would require too many buttons to do so.

This constraint affects object orientation in the same way that Track To does, except this constraint bases the orientation of its poles on the original orientation of the bone! See the page on Tracking for more info. Locked Track also works with this constraint.

Settings

R
: Pressing the R button calculates the rest length as the distance from the centers of the constrained object and its target

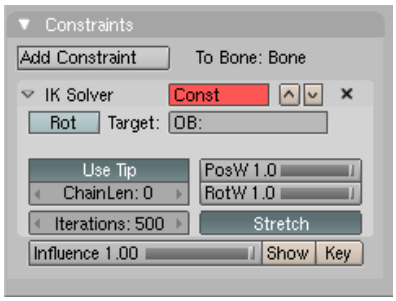
Rest Length
: Rest Length determines the size of the object in the rest position

Volume Variation
: Volume Variation controls the magnitude of the effect

Vol
: The Vol: buttons control along what axes the volume is preserved (if at all)

Plane
: The Plane buttons define which local orientation should be maintained while tracking the target

IK Solver

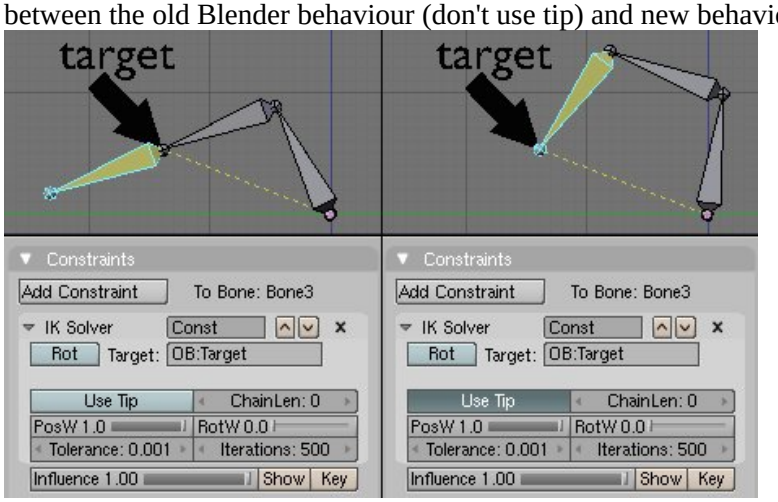


The IK Solver constraint is Blender's implementation of inverse kinematics. You add this constraint to a bone and then it, and the bones above it, become part of the inverse kinematic solving algorithm.

Settings

Rot
: Toggle option to make the IK chain follow the rotation of the target object.

Use Tip
: Toggle option to use the tip of the bone instead of the base to solve the IK to. This option toggles between the old Blender behaviour (don't use tip) and new behavior (use tip).



An IK constraint on the yellow bone with indicated target. Left: without Use Tip. Right: with Use Tip.

ChainLen
: The number of bones above this bone that you want to be affected for IK. The default is 0, which means *all* bones above this bone will be used for IK.

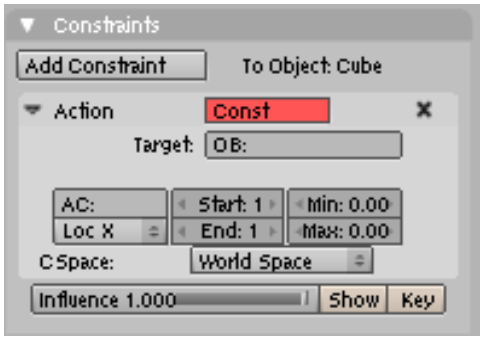
PosW
: Weight of position control for this target.

RotW
: Weight of orientation control for this target.

Tolerance
: foo

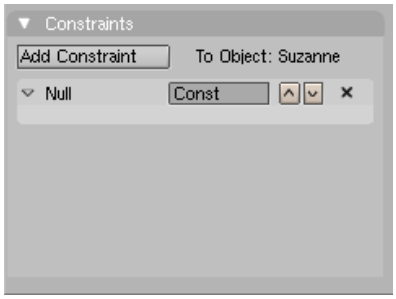
Iterations
: Maximum number of solving iterations.

Action



The Action constraint allows you to map any action to one of the rotation axes of a bone.

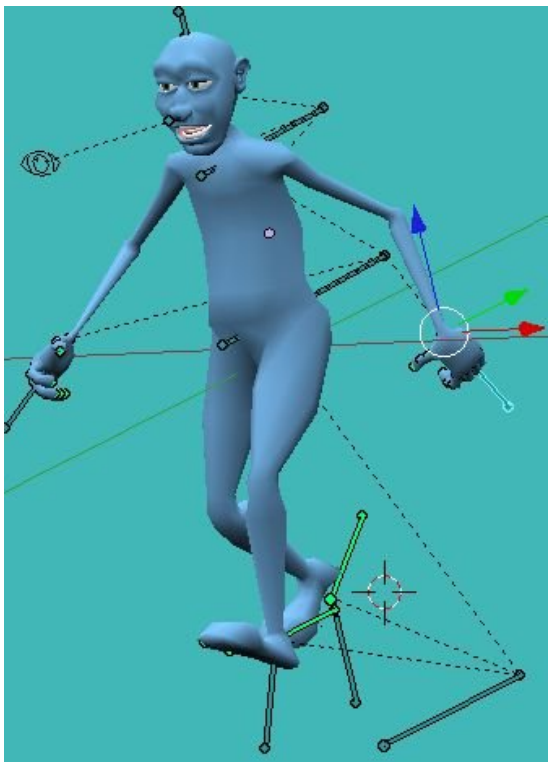
Null



The null constraint doesn't do anything. It is an antiquated feature.

Character Animation using the Armature System

As we have seen in Your First Animation in 30 + 30 Minutes Part II Blender uses *Armatures* for character animation. An armature, once parented to our character mesh, is just like a skeleton which will let us define a number of *poses* for our character along the timeline of our animation. An armature is made up of an arbitrary number of *bones*. The size, position and orientation of every bone in your armature is up to you, and you will find through this chapter that different situations will require a particular arrangement of bones for your character to work properly.



As you animate your armature you will find that it is better to organize several related poses in something called an *action*, which is more or less the same as in the real world. When we walk, we can imagine ourselves passing through several instantaneous poses as if we were in the frames of a moving picture, the whole process of the walk is an action in the end.

But there are actions and actions. As an animator you will need to acquire the capability of knowing how to split any natural movement or action into several simpler actions that will be easier to deal with. Working with simpler actions commonly saves time and work (and why not: money!) since these actions are usually reusable.

Once you have set-up your first actions you will be able to combine them using Blender's powerful *Non Linear Animation* (or *NLA*) editor, giving your character a living mood and natural manners.

In this chapter we will cover every single detail of Blender's functionalities related to Armatures, Actions and the NLA Editor. Furthermore we will see several armature set-ups that will give you a starting point for your own creations and ideas. Relax and enjoy.

Chapters

The Armature Object
Editing Armatures
Posing Armatures

Inverse Kinematics
Mesh Skin Weighting
The Action Editor
Non Linear Animation

Adding an Armature

Mode: Object Mode / Edit Mode (Armature)

Hotkey: â § ShiftA

Menu: Add → Armature

Description

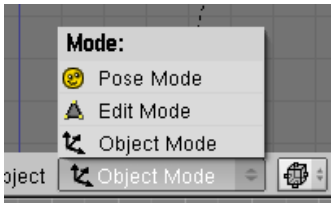
Armatures are like articulated skeletons, that allow you to pose and deform the geometry that surrounds it. An armature is made of a series of bones connected to each other via parenting or constraints. Armatures are most often used in character animation, as a manipulatable skeleton to pose a character, but they can also be useful in other situations too.

An armature is like any other object type:

- It has a center, a position, a rotation and a scale factor.
- It has ObData that can be edited
- It can be linked to other scenes, and the same armature data can be reused on multiple objects.
- All animation you do in object mode is only working on the object, not the armature's contents like bones.

An armature has 3 modes. You can switch using the dropdown menu in the header of the 3Dview or use â Tab to switch between Armature Edit Mode and Object Mode. When in ObjectMode, you can switch the Pose Mode on and off with Ctrlâ Tab. Rather than being an explicit mode that excludes all others, Pose Mode is a state of the armature you can switch on or off. So when in PoseMode, you are still in ObjectMode (you can select another object, contrary to the EditMode).

Options



Mode Pulldown

Object Mode
: Your armature is like any other Object, you can move it around the scene, scale it, rotate it and edit options in the Object button window (F7).

Edit Mode
: Your armature is in what we call rest position, you can move, add and delete the bones it contains.

Pose Mode
: Your armature is ready to be animated, each bone can be moved, scaled or rotated relative to the rest position defined in Edit Mode. Constraints can be applied, you can pose your character, add keys and animate the bone's behavior over time.

Armature Datablock

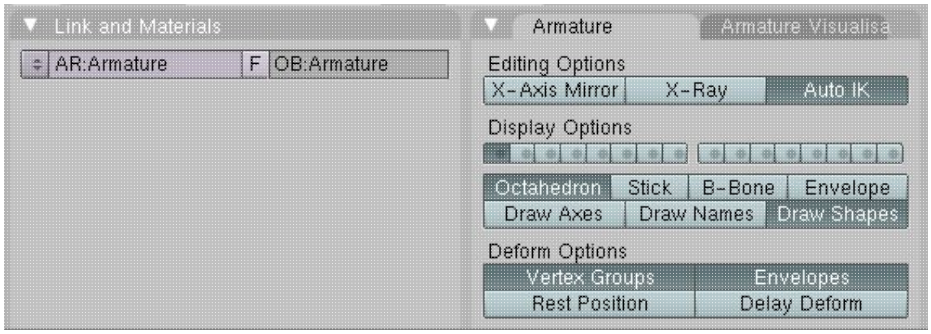
Mode: Object Mode / Edit Mode (Armature)

Panel: Editing Context → Link and Materials

Description

Armatures consist of an Object and an Armature data block. These can be manipulated in the same way as other objects

Options



Armature Panel

AR
: Rename the armature Datablock. The dropdown is a quick way to select which Armature datablock you want to connect to this armature object. You can keep more than one version for the same character. Useful when you have a special move to achieve in a shot, you can turn on an armature for special purposes.

F
: Assign a fake user to the Armature. This is very usefull for when you have more than one armature for your character. An armature that is currently not assigned to a character will not be saved in your .blend file! However, if you assign a fake user, it will always be saved. You can always do batch fake-assignement of armatures by opening the Datablock browser (â § ShiftF4), go back one level (..) to see the root of your project, then go in Armature datablock, select all the armatures you want to keep and press the F.

OB
: Rename your armature object to something more cool and useful than `Armature`, `Armature.001`, etc...

Armature Display Options

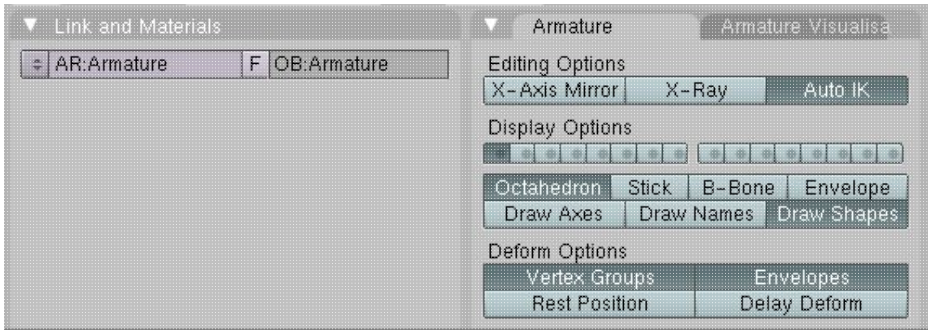
Mode: Object Mode / Edit Mode (Armature)

Panel: Editing Context → Armature

Description

Various methods to control how Armatures are displayed in the 3D View are available. Also note there are some specific options and features that relate to the display mode you're in.

Options



Armature Panel

X-Ray
: The same as in the Object Context, this lets you see the armature through anything in the scene, solid or not. It's useful to see where your bones are in your character so you can select them.

Display Options

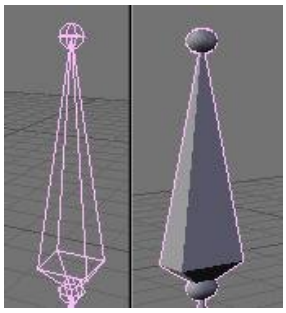
Much like Object Layers, an Armature element (bone, for example) can exist on one or more layers. To work on only those bones of interest to you, select which layers to display and unclutter your view.

Octahedron
: This is the default view. Nothing exciting except you have a good idea of the rolling of the bones.

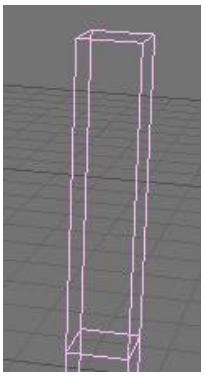
Stick
: This display mode is really useful when you have a lot of bones in your view. It lets you "unclutter" the screen a bit. It draws the bones as tiny sticks.

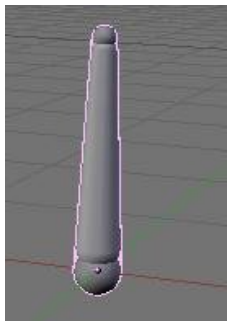
B-Bones
: It's more a feature than a display mode. This is useful to visualise the effect you get when you activate the B-bones (Bezier-Bones). Each join between 2 bones acts like a curve handle and lets you get extremely curvy poses. This is explained further in the Posing Armatures section.

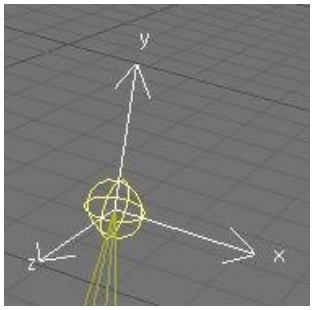
Envelope
: Again it's more a feature than a display mode, using envelopes lets you deform all vertices within a bone's proximity, rather than having to explicitly define and weight paint vertex groups. In this case the visualization will be useful to tweak your rig, showing you which part of the mesh that a bone will move. It's possible interactively change the zone of influence in this display mode. The zone is only visible in EditMode or PoseMode though.


Octahedron


Stick

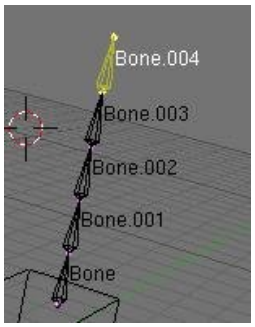

B-Bones


Envelope



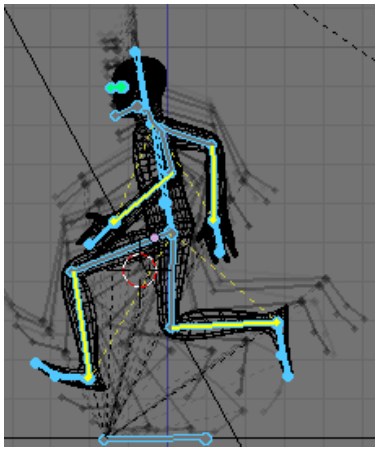
Draw Axes

Draw Axes
: To draw the axes on each bone of the armature when you are in Editmode of PoseMode. Handy when you want to know where you are, and which axis to use in a constraint for example. Mental note: Y is up, Z is depth and X is side, contrary to object for which Z is up, Y is depth and X is side. (*Draw Axes*)



Draw names

Draw names
: This lets you see names of bones whatever the mode you are in. It's useful again to edit your armature, create parent dependencies or add constraints. (*Draw names*)



Ghost

Ghost

Shows a transparent 'ghost' of the armature N frames behind and over the current time. This only works when you have an action linked to the armature, as we will see in the Posing Armatures section.(*Ghost*)

Step

The frame interval between ghost instances.

Armature Deformation

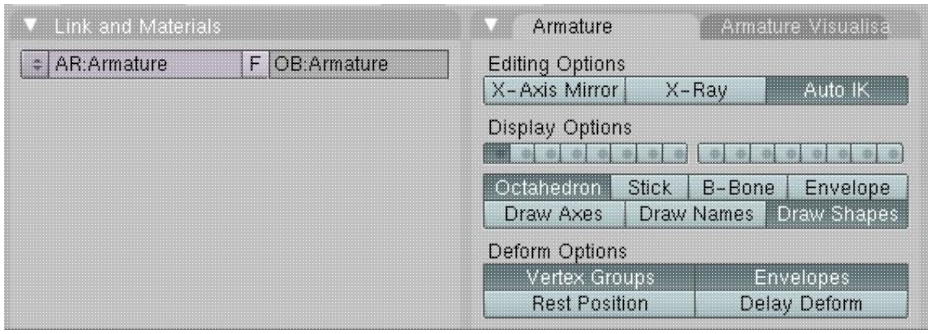
Mode: Object Mode / Edit Mode (Armature)

Panel: Editing Context → Armature

Description

Armatures can deform meshes either through a parent relationship or through modifiers. There are various options to control this.

Options



Armature Panel

Vertex Groups & Envelopes
: Those two toggles let you choose if you want the armature to deform your character using the Vertex Groups and/or the Envelopes. These are used when the Armature is deforming via a parent relationship. If you are using an Armature modifier, these controls are available per modifier in the modifiers panel.

Rest position
: This will show the character as factory default (as defined in EditMode), no actions will be applied to the armature so you can easily edit it in the middle of an animation.

Delay Deform
: This was useful before when the old system was *very* slow. What it does is when you do a manipulation to the rig, it waits until you finish to update the view.

Editing Armatures

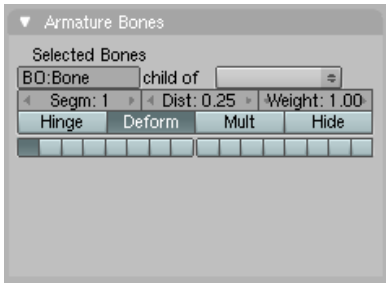
Mode: Edit Mode (Armature)

Panel: Editing Context → Armature

Description

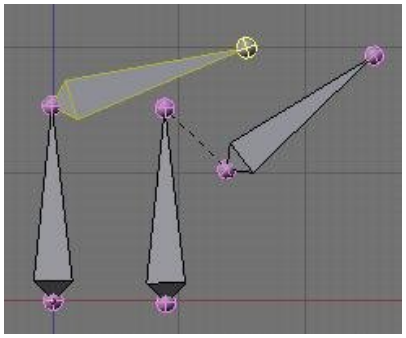
Armatures are comprised of Bones. Editing an Armature in Edit Mode allows you to manipulate the bones in their default rest position.

Options



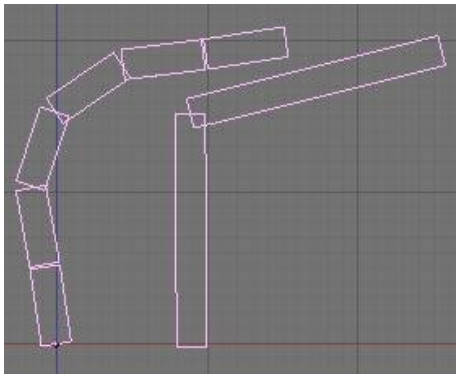
Edit panel: Bones.

BO:
: Rename your bone.



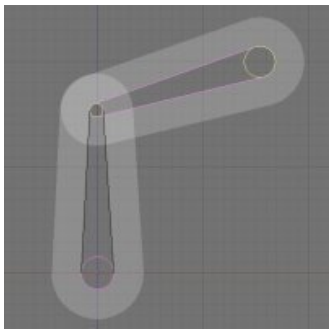
Child of

Child of
: Dropdown: Lets you choose which bone will be the parent of this bone. If so there will be a small button "co" meaning connected. This will set the relationship between your bones. If you parent it to another bone, it will do everything the parent does, rotate, move and scale. A dotted line between the parent and child will appear. If you select Connected, the Root of the Children will stick to the tip of the parent, giving you a chain of bones like the 2 bones in your arm. (*Child of*)



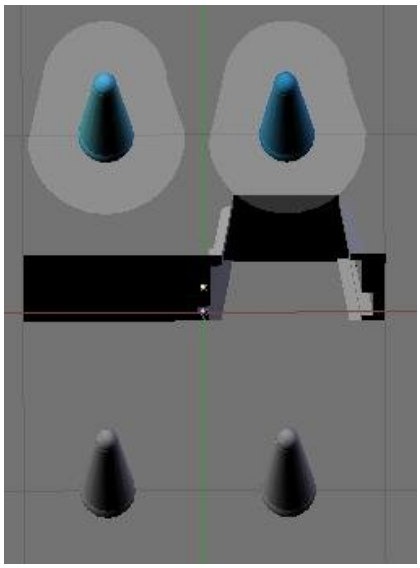
Segmented Bone

Segm
: This is the option to use B-Bones. Set to something > 1, it will cut your bone in many little segments and will deform them on a bezier curve. It really show off in a chain of bones though. See the example between 1 segment (right) and 3 segment each (left). (I returned in Object mode to see the effect). (*Segmented Bone*)



Envelope

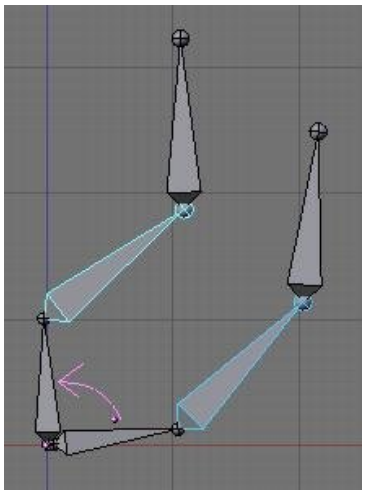
Dist
: This is the area of influence of the bone. it can be visualized using the Envelope display mode. We generally don't touch this field as there is a more easy and fast way to change this option. Turn Envelope drawmode on and select a bone. Then using AltS, you can scale the zone of influence. Better: you can do it on multiple bones and you can do it in EditMode and PoseMode. (*Envelope*)



Weight

Weight

Weight
: This gives you the liberty to tell if this bone will work more or less on the geometry. For example if you have 2 bones working with envelope and crossing each other, you can tell one of them to get more power by lowering the weight of the bone you don't want to use much. If both bones have the same weight (like 1:1) they will influence the geometry equally. But if you set one to 0.5. The other one at 1 will influence more. For example in this image, 2 bones using envelope influence try to move the same geometry. The 2 on the left have the same weight, you can see the geometry didn't move. On the right one of the bones has 0.5 so the bone with 1 of weight is winning the pulling contest!. (*Weight*)



Hinge

Hinge
: This tells the bone to remain motionless in a chain. It doesn't copy the rotation and scale of the parent. Useful for mechanical rig I would say, as you can animate the rotation of the hinge bone without having to correct it because the parent rotated. (*Hinge*)

Deform
: This lets you tell if you want the bone to deform at all. It's like setting weight to 0, except it's faster this way. Useful when using a bone as a target or a controller, i.e. a bone you just want to use to work on other bones.

Mult
: To deform geometry you can use vertex group and/or Envelope. The fact that you can mix both ways gets interesting when you can use one to better tweak the other one. Like using envelope everywhere but tweaking a difficult place manually with vertex group. It's gonna be shown more in details in the Skinning section.

Hide
: This option lets you hide the bone. You can use it to hide useless bones when you try to see what you're doing or just to get a functional rig for an animator when everything is done, kind of "hiding the useless". For example, when you animate you don't need the entire chain of the leg, just the controllers. This option is valid for both EditMode and PoseMode. It's possible to directly Hide bones in the 3Dview by selecting the bones to hide and do H. You can turn all bone visible with AltH too.

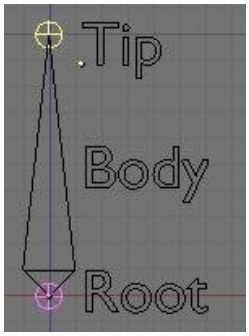
Selecting Bones

Mode: Edit Mode (Armature)

Hotkey: RMB

Menu: Select → Select/Deselect All, Select → Border Select

Description



Bone parts
Bones in Edit mode consist of the tip (the end pointing to) and the root (the end pointing from), and these can be moved independently. RMB clicking on the tip or the root selects either, or clicking on the bone's body selects the entire bone. Selecting both the tip and root selects the entire bone implicitly.

L selects a chain of connected bones under the mouse pointer

P selects the parent bone of selected bones. This is useful when using a fullscreen 3d-View, so an Outliner view is not required to be able to select the parent bone(s). This also works in pose mode.

Adding and Deleting Bones

Mode: Edit Mode (Armature)

Hotkey: â § ShiftA, E, â § ShiftD

Menu: Armature → Extrude, Armature → Duplicate

Description

In Edit mode, you can add a new bone at the cursor location via the Add menu (â § ShiftA).

E extrudes a new bone from a selected root or tip. This will create a bone connected to the original one, meaning the Root of the new bone will follow the Tip of the original one. You can also CtrlLMB to extrude a new bone. It will extrude to where you clicked.

â § ShiftD duplicates selected bones. An entire bone must be selected, not just the root or tip.

X deletes selected bones

Interactively Setting Bone Roll

Mode: Edit Mode (Armature)

Hotkey: CtrlR

Description

Setting the roll value of bones manually has always been a chore. Sometimes, especially on more complex rigs, it may be necessary to manually assign the roll values. However, adjusting the roll value could only be done for one bone at a time, and it was difficult to gauge what the right amount to adjust it was.

Now there is a new transform mode/tool in EditMode. You can select multiple bones and set their roll values interactively. It acts like rotating the bones around the local y-axes in PoseMode. For better feedback, it is recommended that you turn on *Draw Axes* for the Armature in question.

You can access this tool by pressing CtrlR

Subdivide

Mode: Edit Mode (Armature)

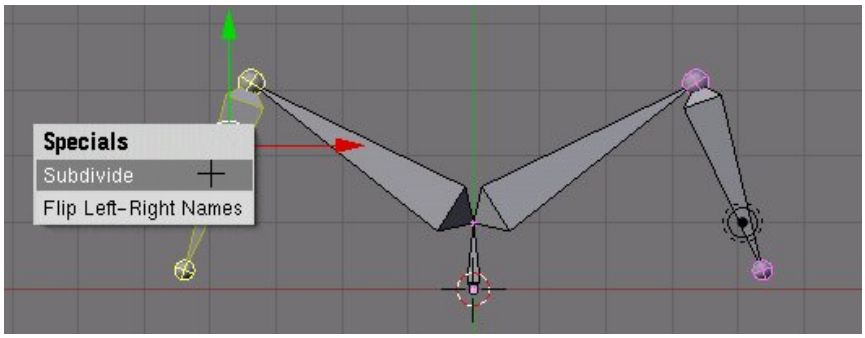
Hotkey: W

Menu: Armature → Subdivide

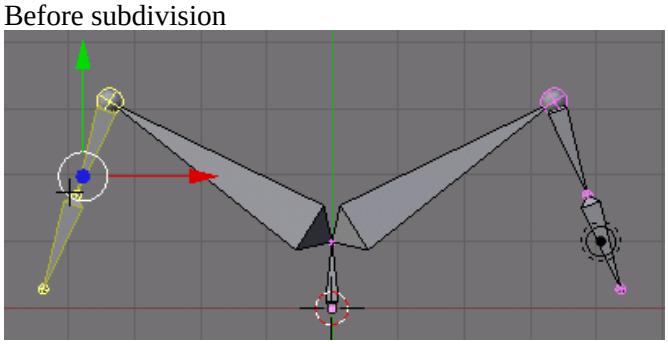
Description

You can subdivide a bone (*Before/After subdivision*) with the W menu. The menu gives you the following options:

- Subdivide - this simply divides the bone into two bones, as shown in the graphic.
- Subdivide Multi - this option gives you a pop-up that asks for the number of cuts. It defaults to the last value used. If you want to change one bone into five, you'll want four cuts. Use the triangles on either side to increment/decrement the field, or click on the field and enter the number you want yourself. Press the OK button when you're done to complete the operation.
- Flip Left-Right Names - Blender will try to help you name bones assuming a right/left symmetry. Sometimes, it gets right and left mixed up, use this option when you know that's going to happen, so you don't have to manually rename the bones.



Before subdivision



After subdivision

X-Axis Mirror Edit

Mode: Edit Mode (Armature)

Panel: Editing Context → Armature

Description

X-axis mirroring replicates edits to one side of an Armature, mirrored on the other side. This works many of Blender's armature editing tools, including manipulations such as move, rotate, scale, for extrude and subdivide. It's a clean way to just do half the job.

The axis of mirroring is X so left<-->right in frontview (1 NumPad) and the center is the center of the armature object.

Flipping Left and Right Bone Names

Mode: Edit Mode (Armature)

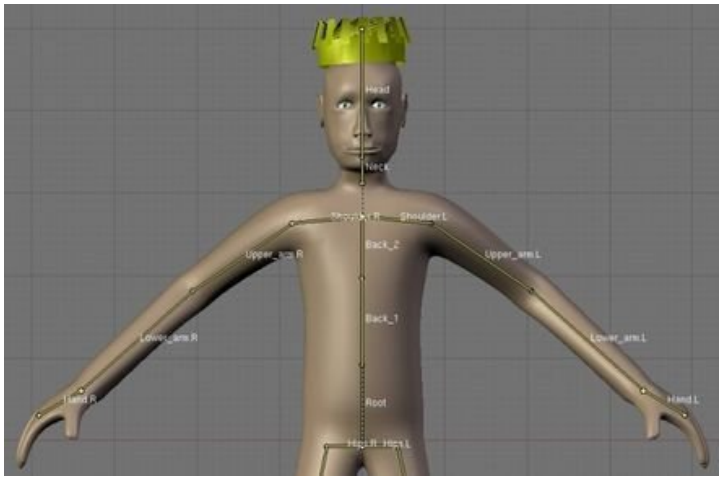
Hotkey: W

Menu: Armature → Flip Left & Right Names

Description

Blender can flip left or right markers in bone names. This can be useful if you've constructed half of a symmetrical rig (marked for a left or right side) and duplicated and mirrored it, and want to update the names for the new side. Blender will swap text in bone names according to the following naming conventions.

Naming conventions



An example of left/right bone naming in a simple rig
Naming conventions in Blender are not only useful for you to find the right bone, but also to tell Blender when any two of them are counterparts.

In case your rig can be mirrored in half (i.e. it's bilaterally symmetrical), it's worthwhile to stick to a left-right naming convention. This will enable you to use some tools that will probably save you time and effort.

- First you should give your bones meaningful base-names. Like `leg`, `arm`, `finger`, `back`, `foot`, etc...
- If you have a bone that has a copy on the other side (a pair), like an arm, give them one of the following separators...
 - Left/right separators can be either the 2nd position (`L_calfbone`) or last-but-one (`calfbone.R`)
 - If there is an lower or upper case `L`, `R`, `left` or `right` blender handles the counter part correctly. For a list of separators see below. Pick one and stick to it as close as possible when rigging, it will pay off. For example:

Lefthand	->	**Right**hand
L Hand.005	->	**R** Hand
hand**.r**	->	hand**.l**
Foot-**l**	->	Foot-**r**
pelvis **LEFT**	->	pelvis **RIGHT**

- Before Blender handles an armature for mirroring or flipping it first removes the number extension, if it's there (like `.001`)
- You can copy a bone named `bla.L` and flip it over using W >> flip name. Blender will name the copy `bla.L.001` and flipping the name will give you `bla.R`. Extensions such as `.001` are also preserved.

Possible separators for Left-Right extensions:

- space " "
- dot "`.`"
- minus "`-`"
- underscore "_"

- If no 'switch' done yet, it tries if a name starts or ends with `left` or `right`, case insensitive. It replaces this, disregarding separator.

Armature Posing Options

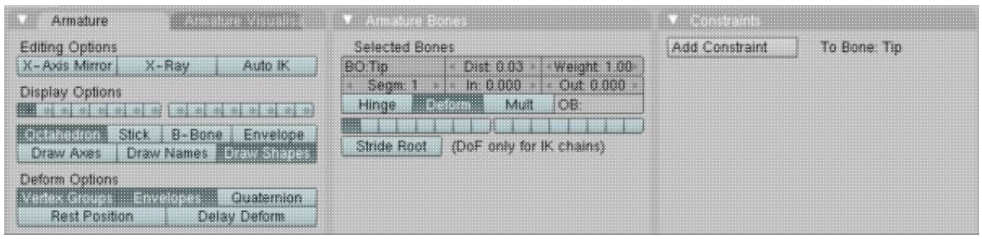
Mode: Edit Mode (Armature)

Panel: Editing Context → Armature

Description

Contrary to EditMode, Pose mode isn't a obligatory mode where you can't do anything else. You can be in posemode and still select another object. When you are done building your armature, you can go into PoseMode to add constraints and start creating actions.

Options



PoseMode panel

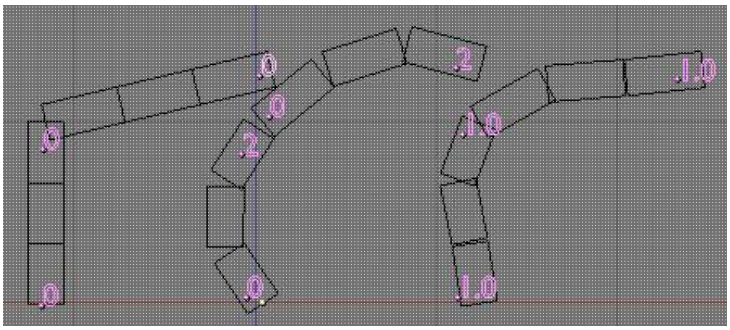
Automatic IK
: Use this feature in the Editbutton (F9) to pose a chain of bones like it was an IK chain. The usefulness is very limited though. It works well only if there is no other IK solver in the chain, and if your chain is isolated from the rest of the rig.

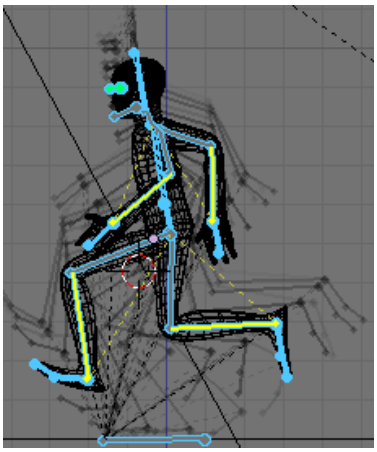
Ghost
: In the Armature Visualisation panel the Ghost option, if set > 0, lets you see the action linked to the armature over time. Also called onion skinning.(*Ghost*)

In / Out
: There are two number fields used to tweak the effect of B-Bones. The In:/Out: is used to tell the scale of the virtual handle of the bezier curve. In: is the root of the bone and Out: is the tip. The bigger the value, the bigger the effect of rotation. (*B-Bones In/Out*)

Limit Rotation

If the bone participates in an IK constraint chain, there will be three more sub-panels on the Armature Bones panel, enabling you to Lock (prevent any changes), set a joint stiffness, or to limit the rotation within a degree min/max. Using these controls, you define an envelope (shown visually) as to the 3D space that the tip of the bone can point to, when following or responding to an IK target.





B-Bones In/Out

Ghost

Armature Posing Tools

Mode: Edit Mode (Armature)

Panel: Editing Context → Armature

Description

- You can pose your rig using G, S and R. Note that if the bone is part of a chain it can't be moved (except if it's the first of the chain, moving all the chain as they are all children), so you rotate the bone instead.
- You can do AltS on one or more bones while in Envelope display mode to tweak the envelope size in real time while animating. Useful when for example you move the hand and some part of the character isn't in the influence zone; the result will be that some vertices will stay behind.
- You can do CtrlC to copy stuff from one bone to a group of bones. The options are location, rotation, scale and constraint. Constraint is very handy when you want to copy a constraint to other bone. The way it works is easy.
- The W menu get some neat options too:
 - ♦ Select constraint target: Will select the target of the bone's constraint currently selected.
 - ♦ Flip name: Yep you can flip name in PoseMode too.
 - ♦ Calculate path/Clear path: This is a visual way to see the action linked to your armature. You can select just some bones and ask Blender to show you the path of the bone.
- You can pose your character and select all bone you want to see included in the action and press I. You can insert a key just for loc, rot or size. Avail will add a key to all available channels in IPO window (all channel you previously added something).
- When you insert a key for your armature, a new action is created and linked to the armature if there was no action before. You can also see the curves of each selected bone of the armature in the IPO window.
- You can parent a bone to an external object by selecting this object then selecting the bone in question so it's active (The armature is in PoseMode so you can select a bone). Do CtrlP. Then when you move the bone the object will follow. This kind of Hard relationship doesn't allow multiple bone influence like a vertice. It's useful when doing robot rigs as you are just moving objects around.

Inverse Kinematics

[Much work has been done on the Inverse Kinematics Solver] both to improve its performance and to add new features. Additionally, setting up IK is now easier than ever!

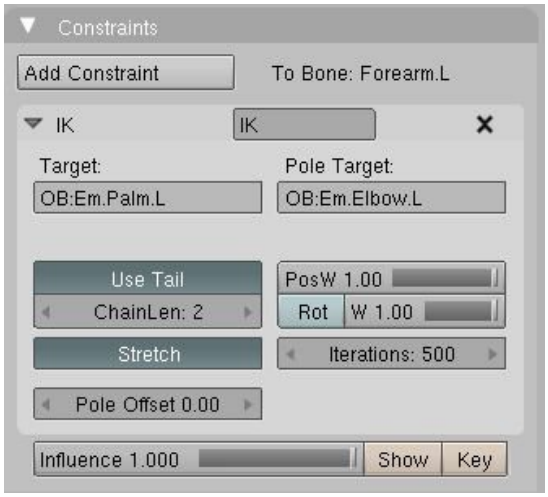
Interface

Mode: Pose Mode

To enable IK solving, you first need something for the bone chain to point to as a target. This target can be an empty, or another bone that is not in the bone chain. There are two ways to initialize the constraint: via hotkey or using the Constraints panel in Pose mode. For this example, select the tip of the bone chain and press â § ShiftI, choosing "To New Empty Object" from the popup menu. A new empty is created, and is set as the target of the IK chain. The bone is colored yellow, and a dashed line indicates how far up the chain of bones the IK solver will work. The new "ChainLen" option of the IK Solver constraint (found in the Object buttons panel) allows control over how far up the chain the IK solver extends, replacing the old "IK" toggle in edit mode. A chain of 0 is **all** bones from that bone back through and connected to it, all the way back to the root bone. This also means that bones in an IK chain no longer have to be connected, but allow an offset from their parent. The other popup menu option, "Without Target", is discussed at the end of this article, under "Mixing IK and FK".

In order to use a bone as an IK target (without creating an empty), select the target bone first, then hold down â § Shift and select the bone at the tip of your IK chain. Pressing â § ShiftI and choosing "To Selected Bone" from the popup menu sets the tip bone to use the target bone for its IK target. Note that the target bone cannot be a part of the bone chain that will have IK.

Tree IK



There is now support for "tree IK". IK chains starting in the same bone will automatically be solved together, affecting each other as needed to reach their targets. The "PosW" slider defines the importance of each target, in case not all targets can be reached.

IK target rotation

The "Rot" option of the IK Solver constraint makes the tip of the chain match the rotation of the target. Its importance relative to reaching the position of the target can be controlled with the "RotW" slider.

IK Pole Target

This bone constraint allows you to define a "pole", which is the direction in which a knee, or middle joint in the IK chain, should point track to. [See Release log for more info.] The bone IK Solver constraint for a bone chain has both an IK target (what the end bone points to) as well as a Pole Target, which is the target that the chain points to, or stays "facing" as it bends. This is most useful in knees and elbows, but would also work to allow you to control which way a tail "crinks" when compressing. Parent an empty to the armature, and then use it as a pole target. Animate the location of that empty, which would guide the bone chain's orientation as it bends. For knees, float the empty out in front of the leg (study mocap to see how the leg/hip rotation results in knee orientation during the walk cycle). For elbows, float the empty behind the arm, and study mocap to see how the elbows fly out or stay tucked in for different movements.



What is a Pole Target?:

"*A pole target is a secondary target for a bone with an IK constraint.*
The first target is where the chain of bones is trying to get to, and the second target (pole target), is where the chain bends to to get to this target.
A possible setup is this: A chain of bones (like upperarm->forearm) has an IK constraint on the last child bone (forearm) and is set to target an IK target (like a replacement target bone for the hand).
Then to control the direction the elbow is pointing, one uses another target, the pole target."

-- Quote (slightly modified) of FreakyDude from this thread.

Rotation limits

The behavior of individual bones in an IK chain can be modified. For bones in an IK chain there are a number of options available in pose mode, in the Edit buttons.

By default bones have 3 degrees of freedom (DoF), meaning they can rotate over the X and Z axes, and roll or twist along the Y axis. Each DoF can now be locked, to disable rotation over that particular axis.

As an example, let us consider a human arm. The wrist has 2 DoF. It can bend in any direction, but it cannot twist. The elbow also has 2 DoF: it can twist, and bend in one direction. Finally, the shoulder has 3 DoF. An example of a 1 DoF joint is the knee.

The "stiffness" defines per DoF how eager the bone is to rotate.

Setting rotation range

The range of rotation can also be limited. "Limit X" and "Limit Z" define how far the bone can rotate over the X and Z axes respectively. If both are enabled this defines an ellipsoid region. "Limit Y" defines how much the bone can twist.

There are two important things to remember:

- DoF and rotation limits are defined with respect to the rest position of the bone.
- They only work for bones in IK chains.

Mixing IK and FK

In many cases you only want to use IK to assist in posing characters, without having an IK defining the motion at all times during an animation. For that purpose, two features were added:

- Targetless IK

When an IK chain has no target defined, it can still be used for posing. Unlike normal IK, you must set keys on all of the bones in the chain if you want to retain the pose. Bones using this type of IK are drawn in a orangish color.

- Automatic IK

This option, in the Editing Buttons, "Armature" Panel, automatically assigns a temporary IK chain to any translated bone, giving the same effect as if the selected bone had been assigned Targetless IK. This chain then only propagates over the connected Bones of the grabbed one.

Skinning

Blender 2.31+
Once the Armature - the 'character skeleton' - is ready it is necessary to parent the character 'skin' to it. Skinning is a technique for creating smooth mesh deformations with an armature. Essentially the skinning is the relationship between the vertices in a mesh and the bones of an armature, and how the transformations of each bone will affect the position of the mesh vertices. When making a child of an armature, several options are presented:

- Parent to Bone

 In this case, a popup menu appears allowing you to choose which bone should be the parent of the child(ren) objects. This is great for robots, whose body parts are *separate* meshes which are not expected to bend and deform when moving.

- Parent to Armature

 Choosing this option will deform the child(ren) mesh(es) according to their vertex groups. If the child meshes don't have any vertex groups, they will be subject to automatic skinning. Indeed a second menu appears, asking:

 ◊ Don't create groups - does nothing else, automatic skinning is used;
 ◊ Name Groups - creates empty vertex groups whose names matches the bone names, but no vertices are assigned to them;
 ◊ Create from closest bone - you want to create *and* populate automatically vertex groups.
 ◊ Create from bone heat - creates vertex groups and assigns weights to each vertex using the 'bone heat' algorithm. This sometimes 'fails to find a solution'; if this happens, in edit mode use P > All loose parts, which will split the mesh into multiple meshes, and then attach the armature to each mesh.

- Parent to Armature Object

 Choosing this option will cause the child(ren) to consider the armature to be an Empty for all intents and purposes.

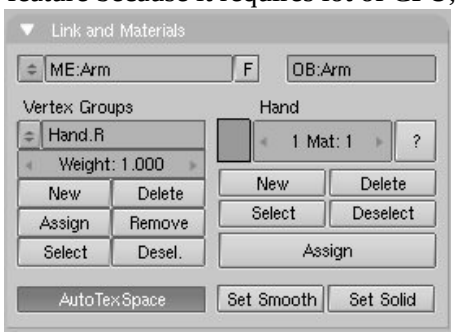
If you are going for character animation then most of the times you will parent your character to the Armature using the "Armature" Option. You are strongly advised to use the Name Groups option. This will provide you with the groups already created, saving the tedious operations of creating an naming them, and possibly avoiding typing errors. The Create from closest bone feature is currently under heavy development. It will use the "Bone types" which can be defined via the menu right of the IK Tog Buttons (*EditButtons for an Armature*) for optimal result. Currently only the Skinnable and Unskinnable options are working. The first option makes Vertex Group be created (and populated, if this is asked for) for the given bone, the second option causes that bone to be ignored in the skinning process.

Automatic Vertex Assignment
The current vertex assignment algorithm creates non-optimal vertex groups, hence it is highly recommended to check each group, one by one.

If a mesh does not have any vertex groups, and it is made the child of an armature, Blender will attempt to calculate deformation information on the fly. This is very slow and is not recommended. It is advisable to create and use vertex groups instead.

Weight and Dist

The *Weight* and *Dist* settings next to the *IK* are only used by the automatic skinning which is a deprecated feature because it requires lot of CPU, produces slow downs and worse result than other methods.



Vertex Groups.

Vertex groups are necessary to define which bones deform which vertices. A vertex can be a member of several groups, in which case its deformation will be a weighted average of the deformations of the bones it is assigned to. In this way it is possible to create smooth joints.

To add a new vertex group to a mesh, you must be in Edit Mode. Create a new vertex group by clicking on the New button in the mesh's Edit Buttons Link and Materials Panel, Vertex Groups group.

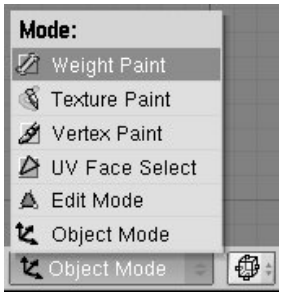
A vertex group can be subsequently deleted by clicking on the Delete button. Change the active group by choosing one from the pull-down group menu. Vertex groups must have the *same* names as the bones that will manipulate them. Both spelling and capitalization matter. This is why automatic name creation is so useful!

Rename a vertex group by Shift-lmb on the name button and typing a new name. Note that vertex group names must be unique within a given mesh.

Vertices can be assigned to the active group by selecting them and clicking the Assign button. Depending on the setting of the Weight button, the vertices will receive more or less influence from the bone. This weighting is only important for vertices that are members of more than one bone. The weight setting is not an absolute value; rather it is a relative one. For each vertex, the system calculates the sum of the weights of all of the bones that affect the vertex. The transformations of each bone are then divided by this amount meaning that each vertex always receives exactly 100% deformation.

Assigning 0 weight to a vertex will effectively remove it from the active group. To remove vertices from the current group select them and click the Remove button. Pressing the Select button will add the vertices assigned to the current group to the selection set. Pressing the Deselect button will remove the vertices assigned to the current group from the selection set. This is handy to check which vertices are in which group.

Weight Painting



Weight Paint Button.

Weight painting is an alternate technique for assigning weights to vertices in vertex groups. The user can "paint" weights onto the model and see the results in real-time. This makes smooth joints easier to achieve. To activate weight-painting mode, select a mesh with vertex groups and click on the weight paint icon (*Weight*

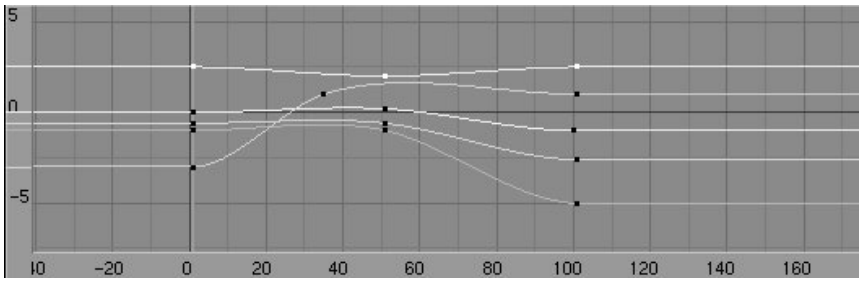
Paint Button.).

The active mesh will be displayed in Weight-Colour mode. In this mode dark blue represents areas with no weight from the current group and red represent areas with full weight. Only one group can be visualized at a time. Changing the active vertex group in the Edit Buttons will change the weight painting display.

Weights are painted onto the mesh using techniques similar to those used for vertex painting, with a few exceptions. The "colour" is the weight value specified in the mesh's Edit Buttons. The opacity slider in the vertex paint Buttons is used to modulate the weight. To erase weight from vertices, set the weight to "0" and start painting.

Baking Actions
If you have an animation that involves constraints and you would like to use it in the game engine (which does not evaluate constraints, and is not covered in this Book), you can bake the Action by pressing the BAKE button in the Action Window ToolBar. This will create a new Action in which every frame is a KeyFrame. This Action can be played in the game engine and should display correctly with all constraints removed. For best results, make sure that all constraint targets are located within the same armature.

You can actually see the Action Ipo associated to a bone in the Ipo Window instead of in the Action Window if you switch to an Ipo Window (*Action Ipo*). The Action Ipo is a special Ipo type that is only applicable to bones. Instead of using Euler angles to encode rotation, Action Ipos use quaternions, which provide better interpolation between Poses.



Action Ipo.
Quaternions use a four-component vector. It is generally difficult and unintuitive to describe the relationships of these quaternion channels to the resulting orientation, but it is often not necessary. It is best to generate quaternion KeyFrames by manipulating the bones directly, only editing the specific curves to adjust lead-in and lead-out transitions.

Context menu

Weight paint mode also has a 'Paint' context menu with the following options:

- Apply bone envelopes to vertex groups - calculates vertex groups and weights from the bone envelopes
- Apply bone heat weights to vertex groups - similar to the option when initially parenting (see above), uses the 'bone heat' algorithm to create vertex groups and assign weights to vertexes
- Various scripts

The Action Editor

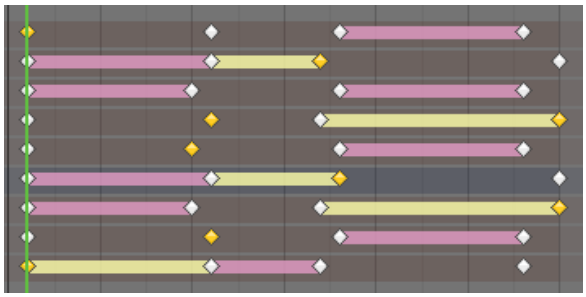
The Action Editor provides control over the keys set for bones in an armature, as well as for Shape Keys.

Long Keyframes

When animating, it is often useful to be able to visually see where the 'pauses' are between keyframes. Long keyframes do this - linking two keyframes in the same channel together.

Long keyframes are only drawn when the two keyframes have the exact same values. This has to happen for every ipo-curve represented by the keyframes shown for a long keyframe to be drawn.

There are two new theme colours for the action editor. They are for the selected and deselected colours of the long keyframes (currently defaulted to be the same as the NLA strip selection colours).



Long Keyframes

Shape Key Sliders

Due to a few internal cleanups of code, it is now possible to save the visibility status of Shape Key sliders. Also, there are no more problems with having multiple Action Editors open, and all displaying Shape-Key actions.

Action Window

An Action is made of one or more Action channels. Each channel corresponds to one of the bones in the armature, and each channel has an Action IPO associated with it. The Action Window provides a means to visualize and Edit all of the IPOs associated with the Action together.

Selection Tools

Selection tools in the Action Editor have always been limited to only a few basic ones. This release, several new options have been added.

Invert Keys
: Selects the keyframes that were not selected, and deselects the keyframes that were selected.

Column Select â On Selected Keys

Selects all the keyframes that occur on the same frame as selected keyframes. Use the hotkey K.

Column Select â On Selected Markers
Selects all the keyframes that occur on the same frame as selected markers. Use the hotkey â § ShiftK.

Column Select â Between Selected Markers
Selects all the keyframes that occur between and on the first and last selected markers. Use the hotkey CtrlK.

Mirror Tools

Like in the Ipo Editor, it is now possible to 'mirror' keyframes over a line. This makes it easier to reverse an action. Four options are currently available:

Mirror Over Vertical Axis
Mirrors selected keyframes using frame == 0 as the mirror-line.

Mirror Over Current Frame
Mirror selected keyframes using current frame as the mirror-line.

Mirror Over Horizontal Axis
Mirrors selected keyframes using value == 0 as the mirror-line.

Mirror Over Selected Marker
Mirrors selected keyframes using the frame of the first (chronologically) selected marker as the mirror-line.

Hotkey for Mirror tools is â § ShiftM.

Snap Tools

Shift-S Snapping

So far, snapping tools have been restricted to 'Snap to Nearest Frame' and only for action channels. Obviously, this is too limited. Now, you can also snap keyframes to the current frame:

- Snap To Nearest Frame
- Snap To Current Frame
- Snap To Nearest Marker

Hotkey for Snap tools is â § ShiftS

It is also worth mentioning, that these now work correctly with actions scaled in the NLA editor. Previously, there could be unpredictable results as there was no correction applied.

Auto-Snapping During Transforms

Many people have requested such an option, as it reduces strain on fingers. There's a new selection-box on the header of the action editor, which sets the mode of auto-snapping for transforms. By default auto-snapping is off.

There are 3 modes of auto-snap:

- Off - transforms per normal
- Frame Step - grid-step transform (may have errors with scaled actions)
- Nearest Frame - true snap-to-frame (takes into account nla-scaling)

These translate to the following hotkeys when transforming:

- Off - no keys press/held (as it's always been)
- Frame Step - Ctrl (as it's always been)
- Nearest Frame - Shift (replaces old shift-key behaviour which was not useful)

Channel 'Protecting'

Action Channels and Constraint Channels can be 'protected' (aka locked). That means that the keyframes for that channel cannot be moved (by any transforms or operations which modify the times the keyframes occur at), be duplicated or destroyed, or have their handle type changed.

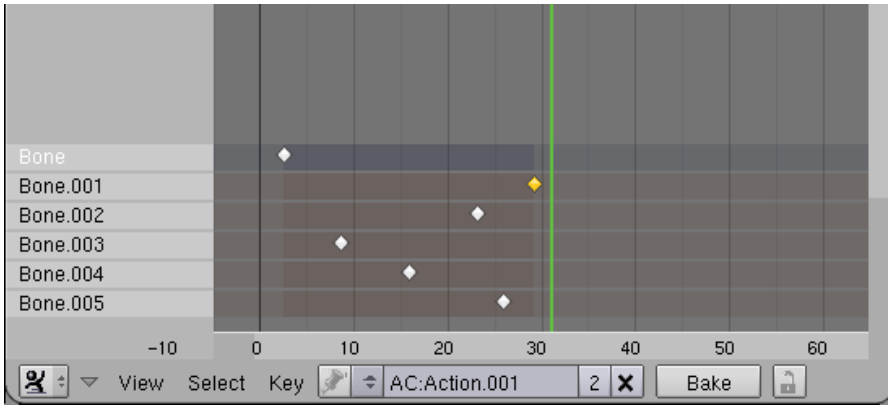
This is useful for protecting parts of animation that has already been finalised, while other parts are worked on.

There is a lock icon on the right-hand side of the channel names. This icon serves two purposes - as an indicator or whether the channel is locked, and also as the way to turning locking on/off for the channel in question.

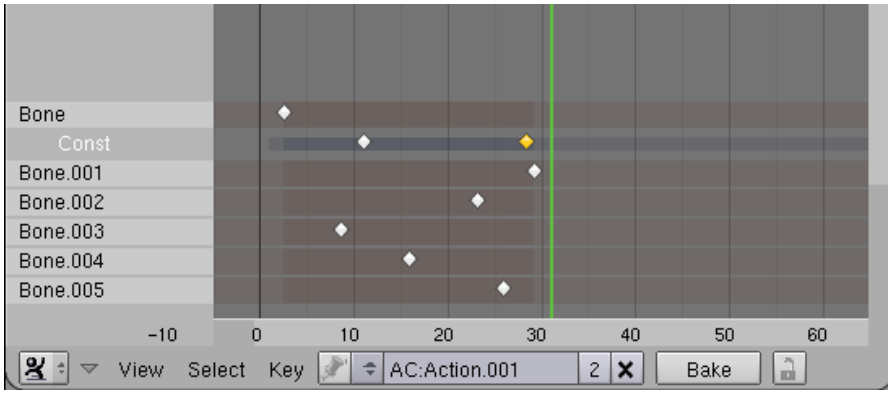
Channel locks also have effect when the action's keys are displayed in the NLA editor. However, they don't currently have any connection with the ipo editor and ipo curves, so it is still possible to insert keyframes and modify keyframes from the ipo editor directly.

Tip

You can activate the Action Window with SHIFT-F12



The Action Window
For every key set in a given Action IPO, a marker will be displayed at the appropriate frame in the Action Window. This is similar to the "Key" mode in the IPO Window. For Action channels with constraint IPOs, there will be one or more additional constraint channels beneath each Action channel. These channels can be selected independently of their owner channels.



Action Window with a Constraint
A block of Action keys can be selected by either RMB on them or by using the boundary select tool B. Selected keys are highlighted in yellow. Once selected, the keys can be moved by pressing G and moving the mouse. Holding Ctrl will lock the movement to whole-frame intervals. LMB will finalize the new location of the keys, while Esc cancels the Action and returns to previous state.

Using AltRMB in the workspace of the Action Editor will select all keyframe markers on that side of the current frame marker.

A block of Action keys can also be scaled horizontally (effectively speeding-up or slowing-down the Action) by selecting number of keys and pressing S. Moving the mouse horizontally will scale the block. LMB will finalize the operation.

Delete one or more selected Action keys by pressing X when the mouse cursor is over the KeyFrame area of the Action Window.

A block of Action keys can be duplicated and moved within the same Action by selecting the desired keys and pressing â § ShiftD. This will immediately enter grab mode so that the new block of keys can be moved. Subsequently LMB will finalize the location of the new keys. Esc will exit grab, but won't remove duplicates.

You can also delete one or more entire Action or constraint channels (and all associated keys) by selecting the channels in the left-most portion of the Action Window (the selected channels will be highlighted in blue, and the names will become white). With the mouse still over the left-hand portion of the window, press X and

confirm the deletion. Note that deleting an Action channel that contains constraint channels will delete those constraint channels as well.

For more detail see the Action Editor Reference

NLA Editor

Mode: Object Mode / Pose Mode (Armature)

Hotkey: Ctrlâ § ShiftF12

Description

You use the NLA (Non Linear Animation) Editor to mix actions and IPOs of objects to make a composite animation. Both Object Ipos and Actions can be blended. Every object may have only one composite animation. You can think of it this way: the Action editor says WHAT the bones are supposed to do, and the NLA Editor says WHEN they are supposed to do it. In addition, you can position the object and key its location using the NLA while in the NLA Editor, in order to say WHERE the action is supposed to occur.

The NLA Editor works in two different modes. In the "Action" (single) mode only the current active action is played. To play combined action use the "NLA" (combined) mode.

In the NLA Editor you see a representation of all Objects with either Actions or Ipos. On a per Object basis, you can add "Action Strips" with â § ShiftA; you can add as many Actions as you like, allowing you to mix different Actions to non-destructively edit timing and relationships. Each Object can have only one active Action, which is displayed in the Action Editor. This active Action is the one which will receive any new keys you insert, and whose keys you can directly edit.

When strips are inserted, they start at the current frame. When a strip is added, it becomes the active strip. The current active strip is always drawn in yellow. Clicking on a strip itself or on the left-hand bars will make the strip active. Using the sript properties panel, you can control how long the action takes to execute, or what portion of an action is to play.

You add new Ipo-Keys to the active Action, you can directly edit these Keys (delete, move). The objects you select in the NLA window are selected also in the 3D window. Be careful! Deleting keys in the NLA also deletes them from the Action, and from the underlying IPO curves as well.

Ipo's for objects - and also for armature objects - can be converted to Actions in the Ipo window with â § ShiftC (or from the Strip Menu (Convert Action to NLA strip)



De-Link Actions before starting:

Before entering the NLA, ensure that any object you want to control using the NLA is de-linked from its action. Otherwise, that action will want to play on its own. To de-link, in the Action Editor, be sure Fake User is enabled and press the X button on the header. The Action will remain in memory (for you to call in using the NLA) but will not force the object to perform it. Note also that Actions designated for use in the NLA should start at frame 1, for convenience.

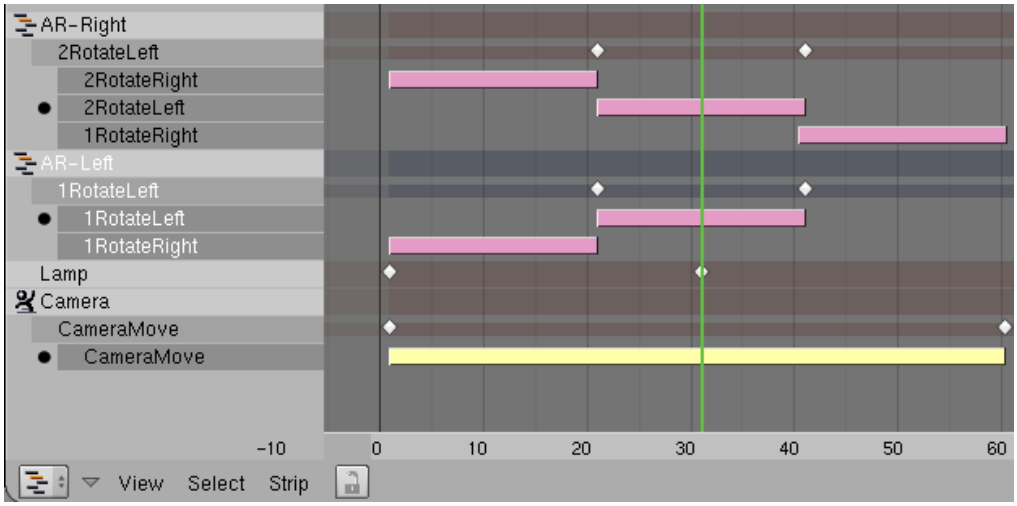


Image 9: Representation of different objects, actions and IPOs in the NLA Editor window.
1) An object is only shown in the NLA editor window if it has Ipos, an action, or NLA strips. The top-most horizontal bar for an individual object shows the Object's name to the left, and any Object Ipo "keyblocks" (the diamonds that represent keyframes) on the right. (*Image 9* - the "Lamp" has only object Ipo's, and no Actions).

2) The button to the immediate left of the Object name toggles Blender between evaluating the entire NLA and displaying the results in the 3D window, or only evaluating the active Action (*Image 9* - "Camera" is in "Action" mode, while "AR-right" and "AR-left" are in "NLA" mode). Note that this icon is only displayed when NLA strips are present for the object.

3) Immediately below the first Object bar, the active Action is shown, and the "keyblocks" for the Action's Ipo's.

4) Finally, the "Action Strips" are drawn. The active Action Strip is indicated with a black dot icon. Clicking on a Strip itself or on the left-hand bars will set a Strip to being the active Action.

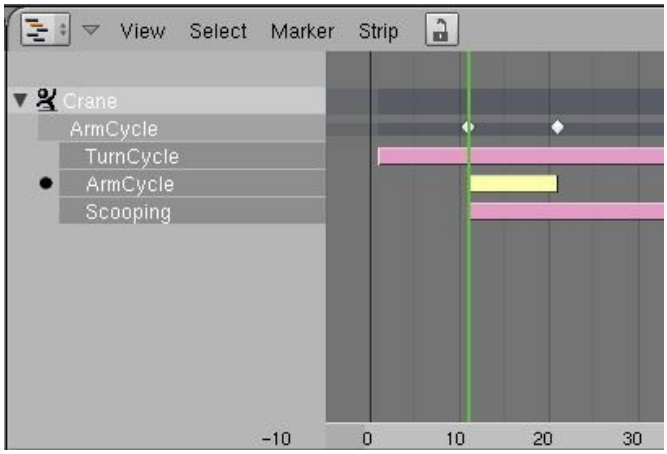
NLA Strip Options

Mode: Object Mode / Pose Mode (Armature)

Hotkey: A

Menu: NLA Editor → Strip → Add Blank Action or Add Action

Description



Strips are added and shown in an indented fashion. when you Add Action Strip, a popup window allows you to choose from available Actions defined. Add Blank Action creates a new action, by default named "ObAction". To rename it, use the Action Window. The name of the overall action is the name of the armature or object. Below that is an indent strip that shows the name of the currently selected strip and all of the ipo keys for that strip. Indented beneath that are all the loaded action strips.

The example image shows an NLA window for armature Crane in Action mode. The RMB selected strip is ArmCycle, and it has two keys. The white diamond at the current frame (green line) shows that there is a keyed position there. The orange diamond at frame 21 is the currently selected frame (it is ready to be Grabbed and moved, or X deleted).

There are two modes for the editor: Action mode and NLA mode. The mode is shown by the icon at the head of the strip: "drowning man" for Action, "mini-strips" for NLA mode. In Action mode, only the action for a single action strip is animated as you scroll through your animation frames. In NLA mode, all strips are animated as you scroll through your animation frames. When you switch to NLA mode, the armature's bones are all recomputed based on keyed IK and FK positions. Any non-keyed bones will revert to their edit position. So, it is a good idea to key all bones (especially IK drivers) at the start of the animation in the default pose.

Recall that to create actions for armatures, you must Add New Action, then Add New IPO strip, or just key the action in the Action window (and an IPO will be added automatically). Also recall that an action can contain motion and IPO curves for many bones. Therefore, it is possible to load two action strips that have two (possibly conflicting) IPO curves for the same bone. For that reason, you have to use the strip's properties to tell Blender which IPO should take precedence, or if it should blend the two actions together. These and other properties are described below.

The selected strip is shown in yellow and has a big dot next to its indent. All keys for that action are shown as diamonds in the indent strip. The window scrolls vertically with LMB -drag the thumb in the right window margin. It scrolls horizontally and vertically via MMB -drag.

Strips are added in order, the last one at the bottom of the list. To move or re-arrange the strips, sue Strip -> Move or the hotkey CtrlPageUp or CtrlPgdn.

Properties

Mode: Object Mode / Pose Mode (Armature)

Hotkey: N

Menu: NLA Editor → Strip → Strip Properties...

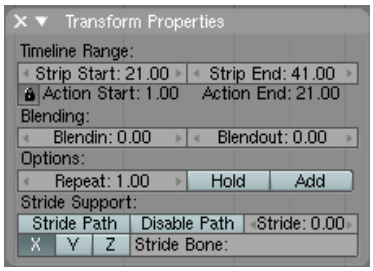


Image 11: The Transform Properties panel for a strip.
The strip properties control how the generic action is applied to the animation. When working on individual actions in Action mode (the "Drowning Man" is shown), the full action is animated in the 3D Window, not just the action between the Action Start and Action End frames in the properties panel. When you switch to NLA mode, only the action range and repeats, along with the rest of the properties, are used to compute the actual animation. These properties include:

Timeline Range
: The first and last frame of the Strip in the timeline. The Strip length is independent from the length of the Action, but the Action will "stretch" to fill the Strip length. This feature allows you to have a generic walk cycle action that spans 40 frames for example. To make your character walk fast, you could enter a 30-frame range; to make him walk slower enter a 50-frame range.

Locked Strip Length
: By setting the strips to Locked, NLA will always keep strip length up-to-date with your Action edits so that all keys are included.
 - ◊ Locked mode is the new default, as it is far superior to the old behavior, but older files will have to be updated manually for this feature to work.
 - ◊ If you release this lock, you can choose only a part of the action to be included in the strip.
 - ◊ Action End is not allowed to be less than Action Start, so you cannot reverse an action in the NLA editor.

Blending
: The number of frames of transition to generate between this Action and the one before it in the Action strip list.

Repeat
: The number of times the Action range should repeat.
 - ◊ This parameter is ignored if Stride Path (in the Stride Support settings) is enabled.

Hold
: If this is enabled, the last frame of the Action will be displayed forever, unless it is overridden by another Action. Otherwise the armature will revert to its rest position.

Add
: Specifies that the transformations in this strip should add to any existing animation data, instead of overwriting it.

Stride Path
: Repeats the strip along a path (for example, to make a walk cycle)

Disable Path
: Temporarily disables the path movement so you can see your action repeating on the spot

Stride
: The distance in Blender Units that the strip should be repeated over (the length of the stride)

X/Y/Z
: The axis of the stride bone to move the armature along

Stride Bone
: The name of the stride bone. See tutorial Stride Support

NLA Strip Tools

Mode: All Modes

Menu: NLA Editor → Strip

Description

Edit the properties of the Strips and Keys via the Strip menu:

Options

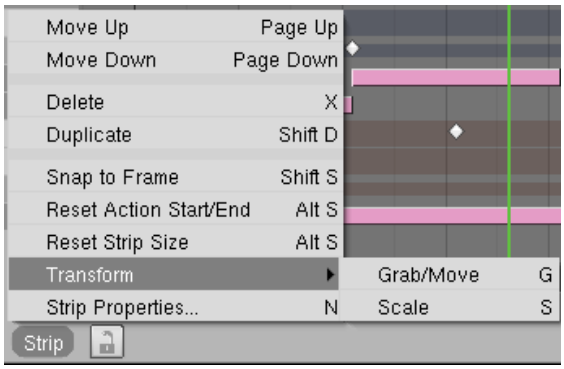


Image 10: The Strip menu.

Move Up/Move Down
: Page up or Page down to change the order of the strips. Strips are evaluated top to bottom. Channels specified in strips later in the list override channels specified in earlier strips.

- Delete

: X Deletes a Strip or a Key. Deleting a Strip affects only the NLA editor. However, if you delete a key the position is ***permanently*** gone.

Duplicate
: â § ShiftD Copies a Strip or an Ipo.

Snap to Frame
: â § ShiftS Snaps start and end of the selected Ipo Strip to frames.

Reset Strip Size
: AltS Resets the size of the Strip in NLA to match the chosen range of frames from Action.

Reset Action Start/End
: AltS Resets the chosen range of Action frames to include all frames in the Action (this is only necessary if you are not using the new "Locked Strip Length" feature, and is useful for fixing older files or if you have "lost your place" when adding or subtracting frames from Actions).

Grab/Move
: G Moves the strip horizontally. With Ctrl you move by frames.

Scale
: S The base "point" for scaling is the current frame. By setting the frame you can select whether the end, the beginning, or a point in the middle of the strip shall keep its position.

Constraints are object features that define spatial relationships between objects, and are the standard method for controlling characters among all 3D animation packages that still implement a more traditional approach to digital character animation. In Blender, constraints can be associated to any type of object or bone object, but not all constraints work with bone objects, and not all constraints work with normal world objects.

Interface

The interface that is used for modifiers and constraints is described here:

- Constraint Stack

Constraints

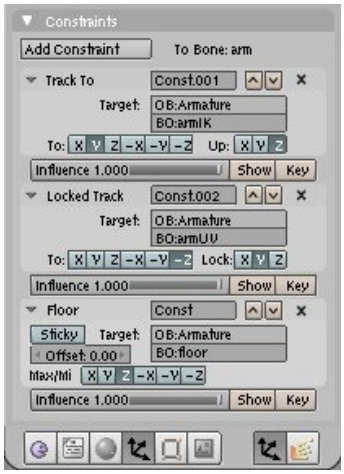
Mode: Any Mode

Panel: Object Context → Constraints

Hotkey: F7 (Panel)

- **Child Of** - Allows a selective application of the effects of parenting to another object.
- **Transformation** -
- **Copy Location** - Copy the position of an object based on some other object's position, so that the objects move together.
- **Copy Rotation** - Copy the rotation of another object so they rotate together.
- **Copy Scale** - Copy the scale of another object.
- **Limit Location** - Object's location is limited to given range.
- **Limit Rotation** - Object's rotation is limited to given range.
- **Limit Scale** - Object's scale is limited to given range.
- **Track To** - Object is tracked to given target object.
- **Floor** -
- **Locked Track** - Object turns to follow a path.
- **Follow Path** - Object moves along a path.
- **Clamp To** -
- **Stretch To** -
- **Rigid Body Joint** -
- **IK Solver** - Bone chain moves to follow another object.
- **Action** - Object executes action based on driving object.
- **Script** - Custom Python script is used as an constraint.
- **Null** -

See Doc:Tutorials/Animation/Armatures/BSoD/Constraints and Axis Locks for descriptions of most available constraints.



Constraints are accessed via the object buttons (F7) in the Constraints panel. After you press the Add Constraint button and select the desired constraint type from the menu, a constraint UI is added to the panel. The constraint will be linked to the active object or the active selected bone, indicated by the label To Object: or To Bone: on the right of the Add Constraint menu.

Constraints are evaluated in a specific order: from top to bottom of the constraint stack. This order can be viewed and changed inside the Constraints panel. Each constraint has a pair of arrow buttons on the right of the constraint name in its top right corner, which can be used to move the constraint up or down the constraint stack. The little cross on the right of these buttons can be used to delete the constraint from the stack.

NOTE: The name field of a newly added constraint will appear red, which indicates that the constraint is not functional. In the case of a newly added constraint this is because the constraint does not yet have a target specified. All constraints require a target object (a normal world object or an armature bone). You should enter the name of the desired target object in the Target: OB field. When you want an armature bone as target, enter the name of the armature object as target. A new field with BO: will appear, where you can place the name of the target bone.

The name field of a constraint will also turn red if the settings for the constraint are invalid or if the constraint conflicts with another constraint. For example: A Track To constraint of which the To and Up vector are set to Z.

Influence

The **influence** of a constraint on the actual location/rotation/size of the object can be defined with the Influence value slider. This can be linked to an Ipo Curve, which can be evoked in the Ipo Curve Editor with the button Show on the right of the Influence value slider. The Key button beside it can be used to add a key to the Ipo Curve.

You can key any constraint, and an object can have multiple constraints keyed to exert a different influence at different times.

Multiple Influence Example

Suppose that you want a camera to follow one path for awhile and then drift off and follow another path. You apply a *follow path* constraint to a camera for one circle/curve/path (make sure you have curve path enabled and enter its name), then open up an IPO curve window. In the buttons window (constraint panel) you can either hit I and insert an Influence value by sliding the slider (to vary the force strength or force fall off) or LMB click the "key" button at the bottom of the constraint panel. Once you have inserted the initial key, up arrow or change to the frame where you want the first influence to start falling off. Go to the IPO window and

and CtrlLMB to insert a key. Go forward a few frames, where you want its influence to be zero, and insert another key, and edit it so that its Y value is zero. At that frame, the constraint in the stack will not influence the camera motion at all. Now insert a *second* Follow Path constraint, and enter the name of the *second* path. Now position back a few frames and insert a key to where you want its influence to start (probably where the first constraint starts dropping off). Edit this second influence curve to come up from zero to one over time. You can imagine that the IPO curves for each circle should cross over each other so that when one is at full influence (max 1.0) the other is at zero etc. You might have to play with the curves a bit to get the transition smooth.

Constraints on Bones

Concerning bones with constraints: the color of the bones in the 3D Window indicate what type of constraint is used:

- **Grey**: No constraint.
- **Yellow**: A bone with an IK constraint.
- **Orange**: A bone with an IK constraint but no target.
- **Green**: A bone with any other kind of constraint.
- **Blue**: A bone that is animated with keyframes.
- **Purple**: The Stride Root.

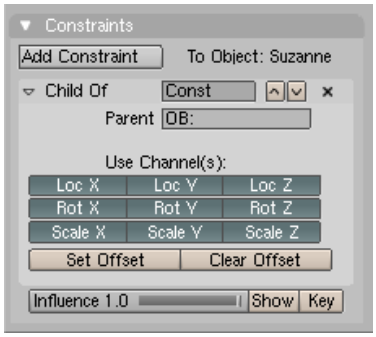
Child Of Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



This certain constraint allows for animatable and or multiple parenting relationships. Now on to the buttons!

Options

Loc X, Y, and Z
: Each of these options will make the parent affect the location according to that axis.

Rot X, Y, and Z
: Makes the parent change the specified axis's rotation.

Scale X, Y, and Z
: Parent affects the scaling on the selected axis.

Set Offset
: Restores the offset (distance) between the objects before the parent relationship.

Clear Offset
: This button moves the object back to where the parent sets it.

Show
: Shows the "Constraints" IPO in the IPO Editor. It will also add a channel for the IPO if there is not already one.

Key
: This little button sets a keyframe for your work all in a simple click.

Tips

When creating a new parent relationship using this constraint, it is usually necessary to click the 'Set Offset' button after assigning the parent. This cancels out any unwanted transform from the parent, so that the owner returns to the position + orientation it was in before the constraint was applied. Note that you should apply 'Set Offset' with all other constraints disabled for a particular child-of constraint.

There are also toggles to enable/disable individual transform channels from the parent affecting the owner. In practice, it is usually best to leave them alone, or disable all of the toggles for a particular transform.

Example

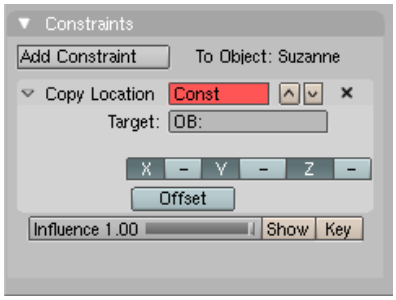
Copy Location Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



Copy Location forces the object to have the same location as its target. When this constraint is used on a bone and another bone is the target, a new button--labeled **Local**--appears next to the **X Y Z** buttons. Using the local button causes the local translation values of the target to be given to the constrained object. In rest position, all bones are at the local location of (0,0,0).

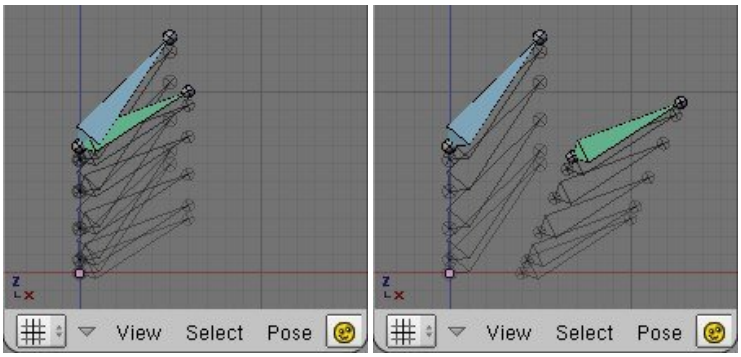
Options

Target
: Object of which location to copy.

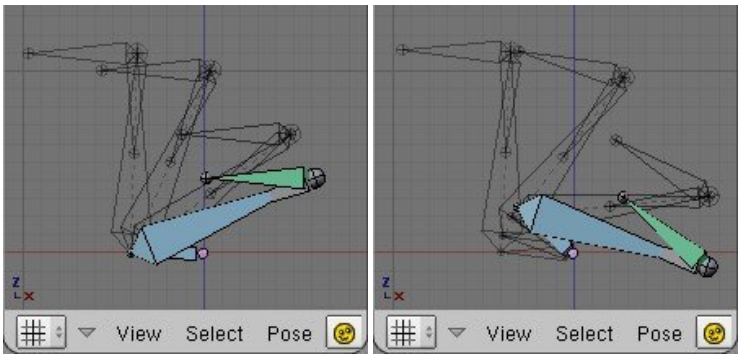
X Y Z
: Axis to constraint. Use - to invert it.

Head/Tail - With Bone targets only
: A number from 0.0 to 1.0 that represents the place on the target bone to use for the Copy (0.0 = the bone's root; 1.0 = the bone's head)

Example



Left: Using global space. This is the default behavior; it's what happens when the local button is not activated.
Right: Using local space, with local button activated.



In these last two images, the constrained bone (shown in green) is actually a child of the root bone (the bone at the beginning of the chain). This demo shows possible uses for the location constraint in a rig. Note that the green bone still inherits rotation from the root because the root is its parent. This is by design though; the green bone could be the child of any of these bones, or none of them.

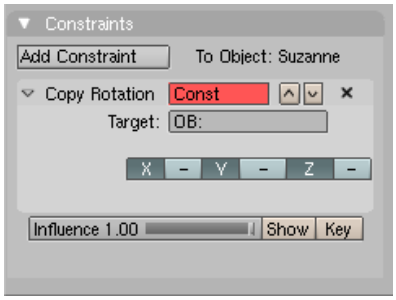
Copy Rotation Constraint

Mode: Object Mode and Pose Mode

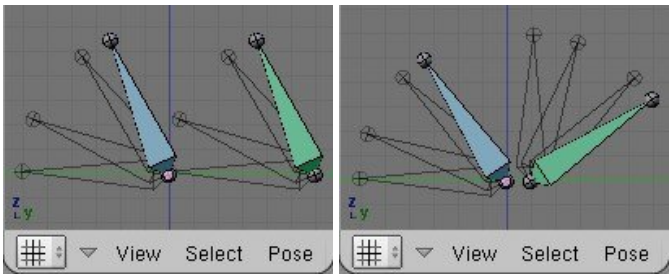
Panel: Object Context → Constraints

Hotkey: F7

Description



Copy rotation causes one object to match the rotation of a target object. For bones, a local option will appear, allowing you to use local space. Imagine you have a bone pointing up and a bone pointing down. If you use local space, each can point different directions, but when the target bone moves to its left, the affected bone will move it to *its* left.



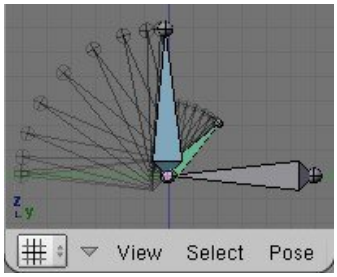
Left: Using global space. Right: Using local space.

Options

Target
: Object of which rotation to copy.

X Y Z
: Axis to constraint. Use the button with the minus sign (-) to invert it.

Example



Here is one good use of the rotation constraint and local space. By setting the influence value to 0.5, the small bone will rotate half as far as the target. This is useful for character joints.

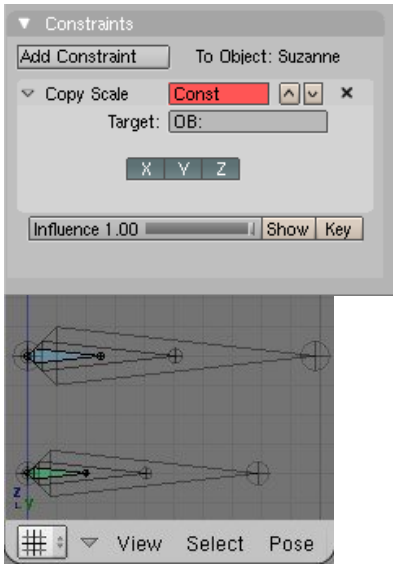
Copy Scale Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



Copy Scale forces the affected object to have the same size as the target. All bones have a (1,1,1) size in rest position. We can draw a bone that is 10 million times longer than the bone right next to it, but in pose mode, they both have a starting size of (1,1,1). You should keep this in mind if you're going to be using the copy scale constraint.

Options

Target
: Object of which scale to copy.

X Y Z
: Axis to constraint.

Example

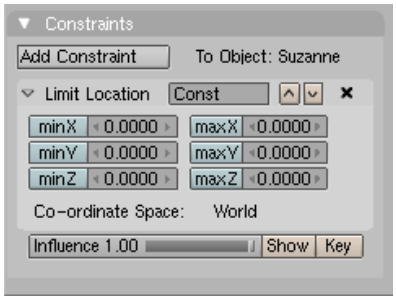
Limit Location Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



An Object can be moved around the scene in the x, y and z coordinates. This constraint places a limit on the location of an object. A limit can be specified on the upper and lower bounds of each coordinate direction. Separate upper and lower limits can be applied for each of the three spatial coordinates (x, y and z). It is interesting to note that even though the limit constrains the visual and rendered location of the object, the object's data block still allows the object to have coordinates outside the minimum and maximum ranges. This can be seen in the Transform Properties N. When an object is grabbed and attempted to be moved outside the limit boundaries, the object will be constrained to those boundaries visually and when rendered but internally, its coordinates will still be changed beyond the limits. If the constraint is removed, the object will seem to jump to its internally specified location. Finally, if an object has an internal location that is beyond the limit, dragging the object back into the limit area will appear to do nothing until the internal coordinates are back within the limit threshold.

The limits for an object are calculated from its center.

Setting a min and max constraint to be the same value constrains the object's movement in that axis effectively locking changes to that axis while allowing motion in the other axis. Although this is possible, using the Transformation Properties axis locking feature is probably easier.

Options

minX, minY, minZ
: These settings specify the minimum location of the object's center in the x, y and z coordinate spaces. To place a limit, the minimum value in world coordinates of the object center location for each axis can be entered. The constraint will not happen unless the associated button for the axis is also pressed.

maxX, maxY, maxZ
: These settings specify the maximum location of the object's center in the x, y and z coordinate spaces. To place a limit, the maximum value in world coordinates of the object center location for each axis can be entered. The constraint will not happen unless the associated button for the axis is also pressed.

Example

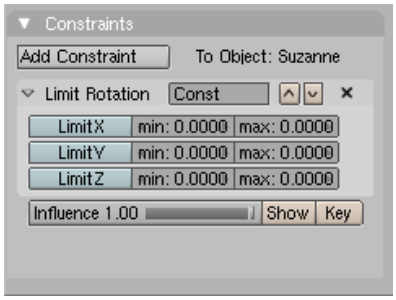
Limit Rotation Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



An Object can be rotated around the x, y and z axis. This constraint places a limit on the amount of rotation. A limit can be specified on the upper and lower values of the rotation around each of the axis. It is interesting to note that even though the limit constrains the visual and rendered rotation of the object, the object's data block still allows the object to have rotation values outside the minimum and maximum ranges. This can be seen in the Transform Properties N. When an object is rotated and attempted to be rotated outside the limit boundaries, the object will be constrained to those boundaries visually and when rendered but internally, its rotation values will still be changed beyond the limits. If the constraint is removed, the object will seem to jump to its internally specified rotation. Finally, if an object has an internal rotation that is beyond the limit, rotating the object back into the limit area will appear to do nothing until the internal rotation values are back within the limit threshold.

Setting a min and max constraint to be the same value constrains the object's rotation in that axis effectively locking changes to that axis while allowing rotation in the other axis. Although this is possible, using the Transformation Properties axis locking feature is probably easier.

Options

LimitX
: The minimum and maximum values of rotation around the x axis. The constraint will not happen unless the associated button for the axis is also pressed.

LimitY
: The minimum and maximum values of rotation around the y axis. The constraint will not happen unless the associated button for the axis is also pressed.

LimitZ
: The minimum and maximum values of rotation around the z axis. The constraint will not happen unless the associated button for the axis is also pressed.

Example

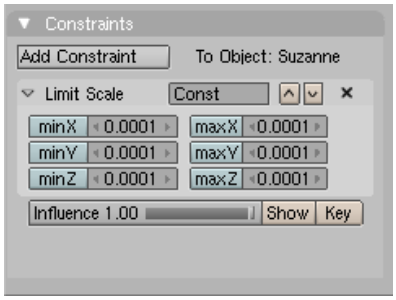
Limit Scale Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



Options

minX, minY, minZ
: foo

maxX, maxY, maxZ
: bar

Example

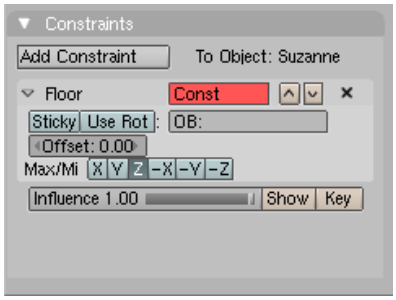
Floor Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



Animation Tip:

When you animate foot placement on the floor plane, ***always*** be sure to use the option **VisualLoc** from the **Insert Key** menu, or enable **Use Visual Keying** from the **Auto Keyframing** menu in the User Preferences.

The Floor Constraint allows you to use a target object to specify the location of a plane which the affected object cannot pass through. In other words, it creates a floor! (or a ceiling, or a wall). This only works by default with planes of the global coordinate system.

Options

Sticky
: Makes the affected object immovable when touching the plane (cannot slide around on the surface of the plane), which is fantastic for making walk and run animations.

UseRot
: Takes the target object's rotation into account.

Offset
: A number of BU's from the object's center. Use this to account for the distance from a foot bone to the surface of the foot's mesh.

Max/Min
: Which axis is the floor? Things normally stick "down" Z, but can walk on walls as well.

Head/Tail - With Bone targets only
: A number from 0.0 to 1.0 that represents the place on the target bone to use for the Floor (0.0 = the bone's root; 1.0 = the bone's head)

Example

Track To Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description

The Track To constraint applies rotations to the object, so that it always points a local â← Toâ← axis towards the target object, with another local â← Upâ← axis permanently maintained as vertical as possible. This tracking is similar to the billboard tracking in 3D. This is the preferred tracking constraint, because it has a more easily controlled constraining mechanism.

The â← billboard trackingâ←
The term â← billboardâ← has a specific meaning in real-time CG programming (i.e. video games!), where it is used for plane objects always facing the camera (they are indeed â← trackersâ← , the camera being their â← targetâ←). Their main usage is as support for tree or mist textures: if they werenâ← t permanently facing the camera, you would often see your trees squeezing to nothing, or your mist turning into a millefeuille paste, which would be funny but not so credibleâ←¦

This constraint shares a close relationship to the IK constraint in some ways. This constraint is very important in rig design, and you should be sure to read and understand the tracking tutorial, as it centers around the use of both of these constraints.

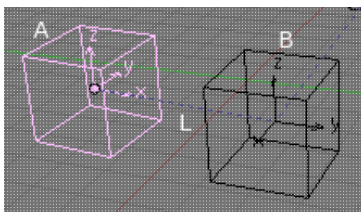
Options

To
: The axis of the object that has to point to the target.

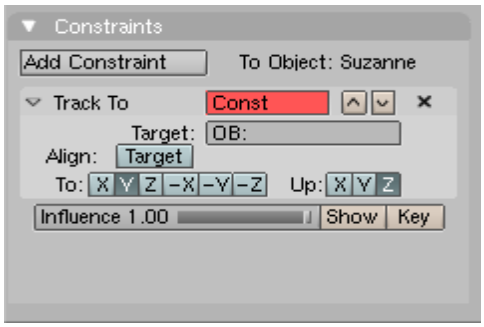
Up
: The axis of the object that has to be aligned (as much as possible) with the world Z axis. An Align: Target button, when enabled, uses the coordinates of the target object's Z axis, thus tilting or rocking the object as it tracks the target.

Head/Tail - With Bone targets only
: A number from 0.0 to 1.0 that represents the place on the target bone to Track to (0.0 = the bone's root; 1.0 = the bone's head)

Example



TrackTo example.

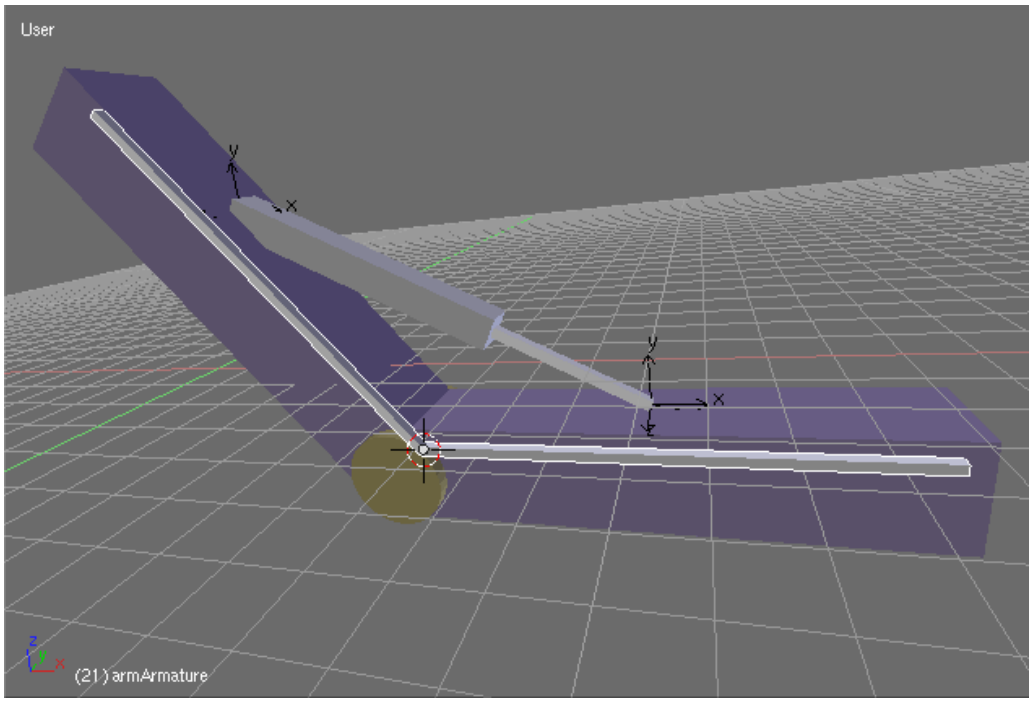


The Constraints panel of the example.

(*TrackTo example*) is an example of a cube using the TrackTo constraint. Cube â← Aâ← is tracking â← Bâ← where â← Lâ← is the tracking line. Notice how the tracking objectâ← s local axes are visible by using the Draw panelâ← s axis button. You can clearly see the tracking â← Toâ← and â← Upâ← axis. Cube â← Aâ← â← s constraint setting are reflected in (*The Constraints panel of the example*). +X is the â← Toâ← axis and Z is the â← Upâ← axis.

You can also see in (*The Constraints panel of the example*) what cube â← Aâ← is tracking by looking at the Target field. We can see that cube â← Aâ← is tracking cube â← Bâ← because cube â← Bâ← â← s name is â← Cubeâ← . You can redirect tracking to another object simply by entering in the name of another object.

Piston



A piston rig created with Track To constraint.
Sample blend file
This constraint is also very useful to rotate the eyes of a character to make the character look at a fixed point. The same can be done for the camera or for lights.

Locked Track Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description

Locked Track is a difficult constraint to explain, both graphically and verbally. The best real-world example would have to be a compass. A compass can rotate to point in the general direction of its target, but it canâ← t point *directly at* the target, because it spins like a wheel on an axle. If a compass is sitting on a table and there is a magnet directly above it, the compass canâ← t point to it. If we move the magnet more to one side of the compass, it still canâ← t point *at* the target, but it can point in the general direction of the target, and still obey its restrictions of the axle.

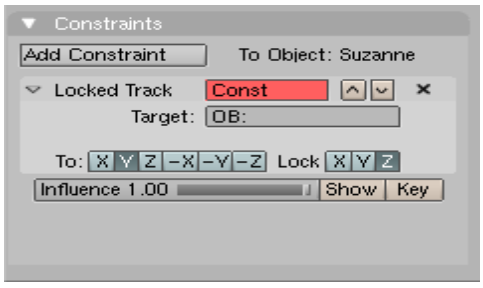
When using a Locked Track constraint, you can think of the target object as a magnet, and the affected object as a compass. The â← Lockâ← axis will function as the axle about which the object spins, and the â← Toâ← axis will function as the compass needle. Which axis does what is up to you! If you have trouble understanding the buttons of this constraint, read the tool-tips; theyâ← re pretty good. If you donâ← t know where your objectâ← s axes are, turn on the Axis button in the Draw panel, Object context (F7). Or, if youâ← re working with bones, turn on the Draw Axes button, Armature panel, Editing context (F9).

This constraint was designed to work cooperatively with the Track To constraint. If you set the axes buttons right for these two constraints, Track To can be used to point the axle at a target object, and Locked Track can spin the object around that axle to a secondary target.

This constraints also works very well for 2D billboarding.

This is all related to the topic discussed at length in the tracking tutorial.

Options



The Locked Track Constraints panel.

Target
: The name of the target object that the tracking object tracks.

VG
: This is only available when tracking to a mesh object. If you type in the name of a vertex group of the target, it will be the center of this vertex group that will be the â← physicalâ← target, not the objectâ← s center.

BO

This is only available when tracking to an armature object. If you type in the name of a bone of the target, it will be this bone that will be the â← physicalâ← target, not the armatureâ← s center.

To
: The tracking axis. It shouldnâ← t be the same as the Lock axis.

Lock
: The locked local axis.

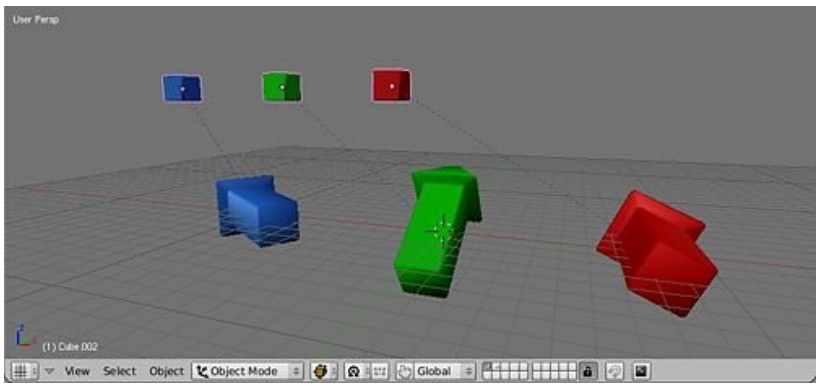
Influence
: This controls how accurately the tracking object tracks the target. **0** means that the constraint is turned off. The tracking object will remain locked in orientation. **1** means tracking is completely on and the tracking axis will stay tightly focused on the target.

Show
: This adds an influence Ipo channel to the constraint if one is not present. You can then add keys to the channel.

Key
: This adds animation keys to the influence Ipo channel. This is a very powerful combination. For example, you could have a camera with a Locked Track constraint applied and have input driving the influence channel.

Example

Compass Arrows



- Sample blend file

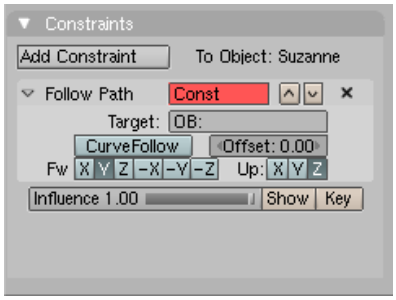
Follow Path Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



Follow Path places the affected object onto a curve object. Curves have an animated property that causes objects along the path to move. In order for this to work, you have to have the CurvePath option activated (it's a property of the curve object). You can find it in the Curve and Surface panel in the Editing buttons, F9.

The movement along the path might be controled by two different ways: the most simple, in this same Curve and Surface panel, is to define the number of frames of the movement *via* the num button Path Len:, and its start frame *via* the constraint's Offset: option (by default: start frame **1** (= offset of **0**), duration **100**).

The second way much more precise and powerful is to define a Speed IPO curve for the path (Path section of the IPO Curve window). The start position along the path will correspond to an IPO value of **0.0**, and the end position, to an IPO value of **1.0**. You can therefore control the start frame, the speed of the movement, and the end frame, and even force your object to go forth and back along the path!

If you don't want objects on the path to move, you can give the path a flat speed IPO curve (its value will control the position of the object along the path).

Follow Path is another constraint that works well with Locked Track. One example is a flying camera on a path. To control the camera's roll angle, you can use a Locked Track and a target object to specify the up direction, as the camera flies along the path.

This constraint does not work well with bones.

Options

Offset
: The number of frames to offset from the "animation" defined by the path (by default: from the frame **1**).

CurveFollow
: If this option isn't activated, the affected object's rotation isn't modified by the curve; otherwise, it's affected depending on the following options:

Fw
: The axis of the object that has to be aligned with the forward direction of the path.

Up
: The axis of the object that has to be aligned (as much as possible) with the world Z axis.
In fact, with this option activated, the behaviour of the affected object shares some properties with the one caused by a Locked Track constraint, with the path as "axle", and the world Z axis as "magnet"

Example

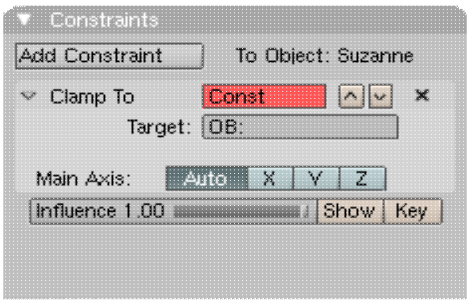
Clamp To Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



Previously, we talked about Follow Path. Now lets talk about the Clamp To constraint.

The Clamp To constraint is a constraint which is particularly useful for moving things along large and complex paths which would otherwise be hard to hand-key smoothly. If you're feeling a little dÃ©jÃ vu, it's because the idea of a Clamp To constraint is very similar to Follow Path, with one main difference.

The difference is this: where Follow Path uses the time IPO of the curve that our constraint is targeting, Clamp To will get the actual position of the object(let's say a switch in some dark and Stygian spiral lock which must be moved back and forth in a combination, instead of the constant paths of the fighter and its pursuers in the example for the Follow Path constraint) and judge where to put the object(switch) by comparing the object's location to the curve it's targeting.

As with most things, of course, there's a bright side and a dark side.

The bright side is that when we're working with Clamp To, it will be easier to see what our object will be doing, since we're working in the 3D view port, in the same window that we're working with, it'll just be a lot more precise than sliding keys around on a time IPO and playing the animation over and over.

The dark(and Stygian) side is that, unlike in the Follow Path constraint, Clamp To doesn't have an option to track our object's rotation (pitch, roll, yaw) to the banking of the targeted curve, but--like our lock, for example, or the armature example further below that--we don't always need rotation on, so in cases like this it's usually a lot handier to fire up a Clamp To, and get the bits of rotation we do need some other way.

All together, what this means is that it'll probably be much easier to animate varied motion across a curve than it might be if we were using Follow Path, but even if it's not easier, it's an interesting alternative, so it can just come down to personal choice.

You don't need to know what Stygian means.

Options

Target

Text box - This is your basic target box, which just displays what your constraint is using as a reference. NOTE: In this case, the target object will have to be a curve object that the constrained object clamps to, so this box only works on curves.

Main Axis

This is a button group at the bottom - selecting one of these picks a global axis(x, y, z) for the constraint to use a reference to how far along the curve it's supposed to be. There's no real wrong choice, so just pick the axis that'd be easiest to work with or works best in your current situation--a good idea is picking the axis which the targeted curve is the longest along, so that the global and constrained locations of the object will be more similar to each other--or you can just pick Auto and Blender will make its best guess.

Example



The Dark and Stygian Lock.

Stygian Lock(object)

.Blend file

This is the Dark and Stygian lock that I've kept going on about, and now you all have to see it. Once again, you don't need to know what Stygian means, but before you all start throwing full wine bottles at your computer screens, looking at the picture on the right should give you a fair idea.

All the fancy design-work aside, though, the face and door are just decoration; the actual animation is done by the Clamp To constraint on the knob. If you watch, you'll see that the motion of it slides around the edge of the disk, goes back and forth, and even stops sometimes, and still manages to keep in a circle. This is the bright side of using a curve. Trying to key this the regular way would probably take much longer, and be much less smooth.

The bright side of, not just using a curve, but actually using a Clamp To constraint in this case--instead of a Follow Path constraint--happens around half-way through the animation. The knob slides all the way around, and then past the face's mouth, and keeps going a little while, before stopping and sliding back. Since the actual curve begins and ends in the figure's mouth, there needs to be an emergency key frame to jump the knob from being near the mouth at the end of the curve, to being near the mouth at the beginning. To do this, we have to visually move the location for both the keys as close to each other as possible, to make the transition smooth, this is something that would be much more difficult with the targeted curve's time IPO.

Note

I decided, in a spur of last-minute inspiration, that the camera would use a Clamp To constraint as well. It may seem a bit belabored, but we might as well go all the way and be consistent. It's a demo file, after all.

Arm(armature)

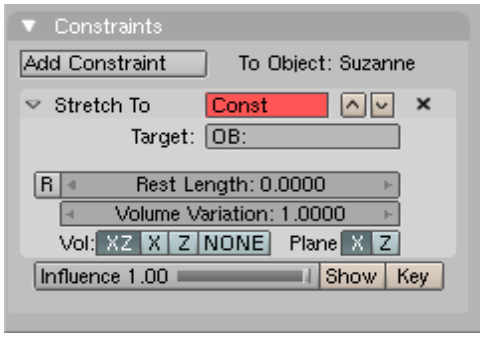
Stretch To Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



Stretch To causes the affected object to scale the Y axis towards a target object. It also has volumetric features, so the affected object can squash down as the target moves closer, or thin out as the target moves farther away. Or you can choose not to make use of this volumetric squash-'n'-stretch feature, by pressing the **NONE** button. This constraint assumes that the Y axis will be the axis that does the stretching, and doesn't give you the option of using a different one because it would require too many buttons to do so.

This constraint affects object orientation in the same way that Track To does, except this constraint bases the orientation of its poles on the original orientation of the bone! See the page on Tracking for more info. Locked Track also works with this constraint.

Options

R
: Pressing the R button calculates the rest length as the distance from the centers of the constrained object and its target

Rest Length
: Rest Length determines the size of the object in the rest position

Volume Variation
: Volume Variation controls the magnitude of the effect

Vol
: The Vol: buttons control along what axes the volume is preserved (if at all)

Plane
: The Plane buttons define which local orientation should be maintained while tracking the target

Example

Transformation Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description

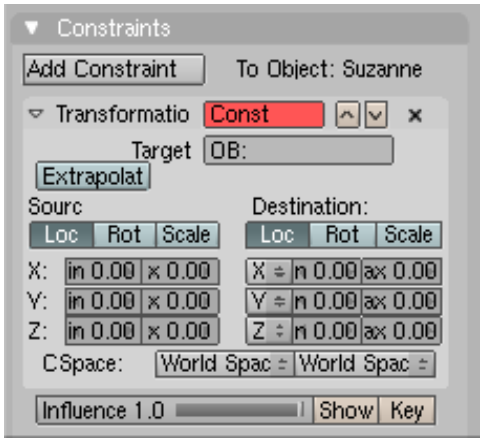
This constraint allows transforms to be converted from one type to another, and also from one range to another. Typical uses for this include gears (*see note), and rotation based on location setups.

NOTE: unfortunately, this will not work for gears, as there are some underlying problems related to the math which mean that this cannot work nicely

The user-interface for this is a bit complex, but can be broken down quite simply:

```
* Left = Source (from Target) Transforms, Right = Destination (for Owner) Transforms
* Controls on Left: Ranges for each axis (x,y,z) of source transform
* Controls on Right: Which source channel (x,y,z) to read from for each channel, and range t
* Loc/Rot/Scale toggles control which transform to affect
* Extrapolate toggle is used to specify whether values outside of ranges should be extrapola
```

It should be noted that when mapping transforms to locations (i.e. Destination = Loc), the owner's existing location is added to the result of evaluating this constraint.



Options

Extrapolation
: description here

Source
: description here

Destination
: description here

CSpace
: description here

Example

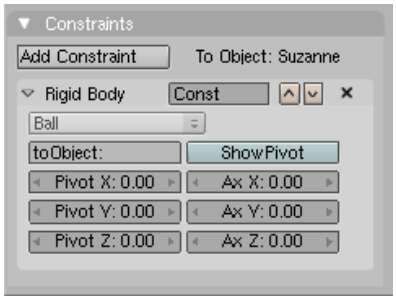
Rigid Body Joint Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



This constraint is mainly useful when using the game engine.

Options

Joint Types
: Ball: works in 2 axis allowing circular movements.
: Hinge: works in one plane, like an elbow.

toObject
: Is the object name that the mesh would be attached to.

ShowPivot
: draws the pivot of the joint on the 3D viewport.

Pivot X, Y, Z
: Relocates the pivot.

Ax X, Y, Z
: Rotates the pivot.

Example

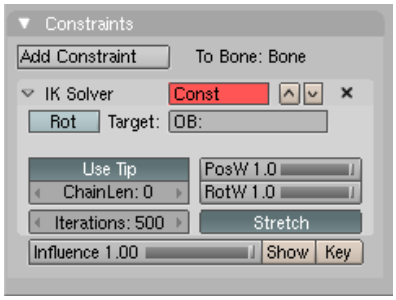
IK Solver Constraint

Mode: Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



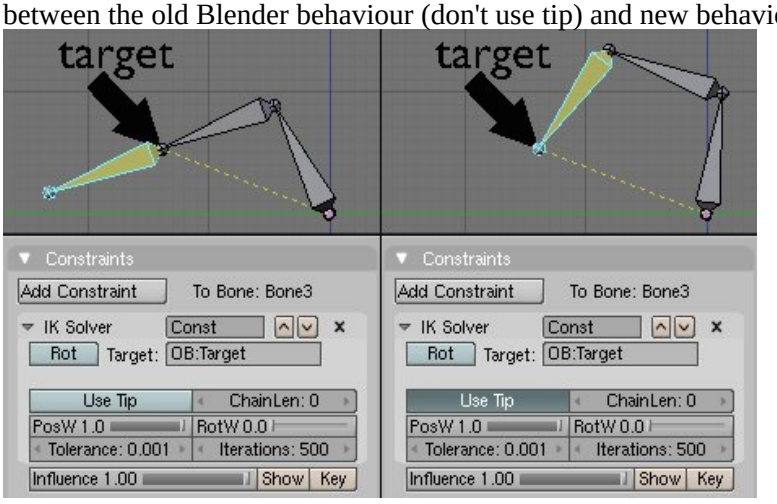
The IK Solver constraint is Blender's implementation of inverse kinematics. You add this constraint to a bone and then it, and the bones above it, become part of the inverse kinematic solving algorithm.

Options

Bone as Target
if you want to use a bone as a target, enter the name of the Armature as the Object name. A Bone field will pop in and allow you to enter the specfic name of the bone to be pointed to.

Rot
: Toggle option to make the IK chain follow the rotation of the target object.

Use Tip
: Toggle option to use the tip of the bone instead of the base to solve the IK to. This option toggles between the old Blender behaviour (don't use tip) and new behavior (use tip).



An IK constraint on the yellow bone with indicated target. Left: without Use Tip. Right: with Use Tip.

ChainLen

The number of bones above this bone that you want to be affected for IK. The default is 0, which means *all* bones above this bone will be used for IK.

PosW
: Weight of position control for this target.

RotW
: Weight of orientation control for this target.

Tolerance
: foo

Iterations
: Maximum number of solving iterations.

Example

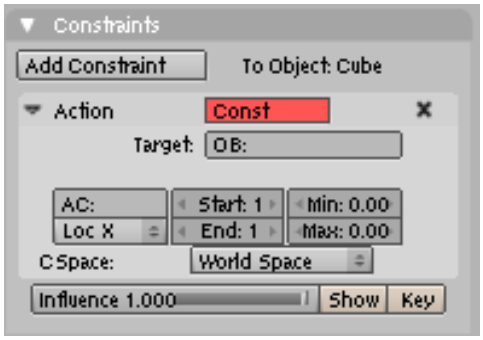
Action Constraint

Mode: Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



The Action constraint allows you to map any action to one of the rotation axes of a bone.

Options

Target
: the object to use in the constraint.

AC
: The action containing the keys for the specified bone.

Loc
: Choose which transformation axis (from the object) is used to set keys.

Start
: Starting frame of the action

End
: Final frame of the action.

Min
: The minimum value for the target channel range.

Max
: The maximum value for the channel range.

C Space
: The space (Global or Local) that the object is evaluated in.

Example

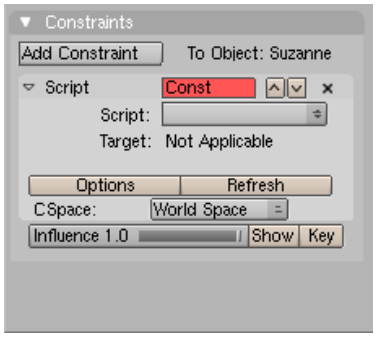
Script Constraint

Mode: Object Mode and Pose Mode

Panel: Object Context → Constraints

Hotkey: F7

Description



The script constraint allows a person to write a new constraint, all in Python. A Pyconstraint can have multiple targets (however many the script requires), and even their own options.

Options

Script
: Sets the Pyconstraint to use.

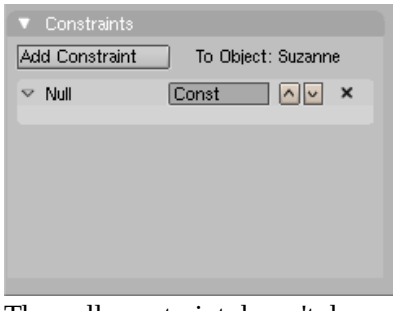
Target
: Allows you to select what objects the script will target.

Options
: Change some of the constraint's settings.

Refresh
: Forces the scripted constraint to refresh it's settings.

C Space
: The space (Global, Local) which the target is used.

Example

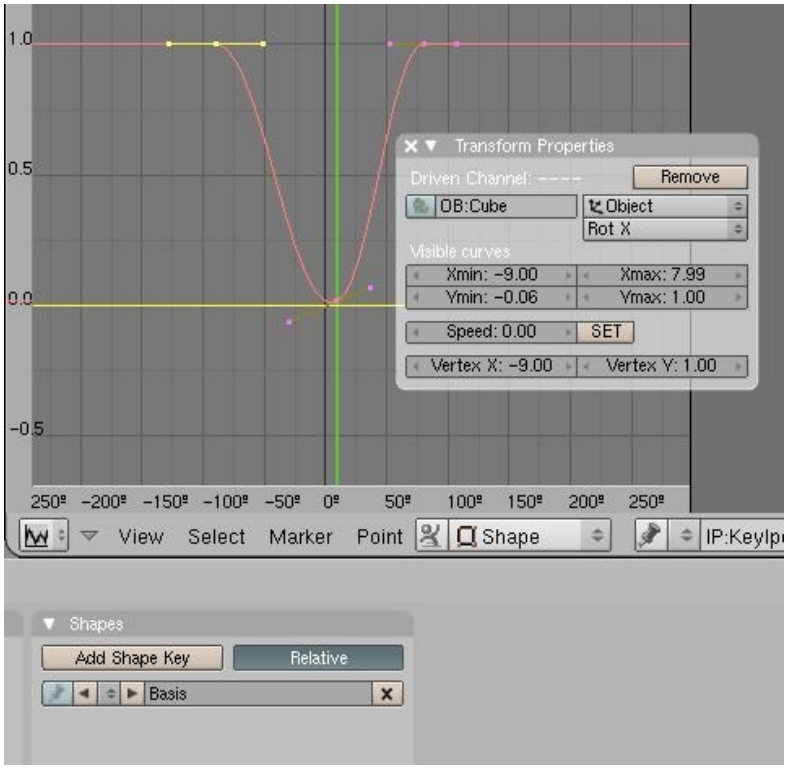


The null constraint doesn't do anything. It is an antiquated feature.

Introduction

A shape key can be made to happen by something happening to another object. Animating that other object causes the shape key to exert its influence on the mesh. This other object can be any kind of object, but is commonly a mesh object that has been modeled to be be a custom shape that indicates what it does. For example, a mesh that controls the eyelids of a face may be a crescent shape.

As an analogy, think of a robot on a platform, with a control panel nearby. As you move a lever on the control panel, the servos in the robot activate and the robot moves its arm. In this example, the driver is the lever, and the shape key is the robot's arm. By moving or rotating or scaling the driver, the shape key activates.



The relationship between the change in the driver's location/rotation/scale and the influence on the shape key is through the Shape's IPO curve. With the base object's shape key selected,

- go to the IPO window and change the type to Shape.
- In the IPO window, press N to bring up the properties panel.
- Click the tan button called *Add Driver*.
- In the OB: field, enter the name of the object that you wish to use to control the influence of the shape key. In this example, the object is called "Cube"
- Select what animation aspect of the driver you wish to use to control the shape key influence. In this example we have used the "Rot X". Depending on what you choose, the units of the IPO window will change to either Blender Units or Degrees.
- CtrlLMB click to add IPO control points. A height (Y) indicates the percent influence that shape key has; 0 is nothing, and 1.0 is 100 percent. The X value indicates what value of the driven channel (Rot X in this example) corresponds to the Y value.

In this example, if Cube is rotated about the X axis by less than -90 degrees or more than 90 degrees, the Basis shape will have 100% influence. For values between -90 and 90, the influence changes from 1 to 0 and back up to 1 again. If the RotX of the cube is 0, the basis shape has no influence. As the cube rotates from 0 to 90, the Basis shape will exert more and more influence.

Now, in your animation, add IPO keys for the Cube, and rotate it about the X axis. As you do, the shape key on the other mesh will activate and deform the mesh. Keep the driver mesh out of camera view, or on a hidden layer, or transparent so that it does not render.

You can see that a single driver object can drive up to 9 different shape keys individually; one for each LocXYZ, RotXYZ and ScaleXYZ

Driven Shape Keys

A *Driven Shape Key* is the controlling of the *Key Value* of a relative *Shape Key* by the movement of another object.

If one would like to animate relative *Shape Keys* in Blender, this unfortunately, cannot be done with **Actions**. (See Action Editor)

Instead one uses an "auxiliary construction", Ipo Drivers. With an Ipo Driver you control the value of an Ipo curve by moving another object. So e.g. the color of an object can be dependent on the rotation of another object. Equally the *Key Value* - the influence strength - of a shape can be dependent on the rotation or movement of another object. This is especially useful, because one can make e.g. the form of a muscle dependent on the amount of rotation of the underlying bone.

Driven Shape Keys are a typical Blender function: built up from several elements, which unfold their power only in the correct combination. In this case that is the vertex animation by relative Shape Key, and the composition of actions in the NLA for Armatures. We do not have to use Armatures (and/or Bones) as "drivers" for the Shape Keys necessarily, but the animation possibilities with Poses are best. Therefore the entire functionality is already described in other sections, it's just that it's not too obvious to combine the elements.

By using driven Shape Keys you're working on increasingly higher levels of abstraction.

- The simplest thing is to work on the level of the Shape Keys, like "Left Eyebrow Down", "Right Eye Open", "Sneer Left".
- The next level is to drive one or more Shapes with the same armature bone, like "Eyelid" + "Eyebrow". Or "Bulge muscle if arm is rotated".
- The next level is to combine the movement of several bones together in an action, like "Happy", "Sad", "Wink", "Frown" ...
- The last level would be to combine the actions in the *NLA Editor*.

You don't have to use these things, but they may make your life and workflow easier.

Setup the Shape Keys

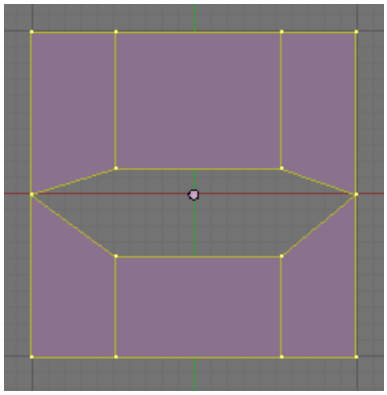


Image 1: Base Mesh for the Shape Keys example.
We will start with a simple example, which shows all essential elements.

The Basis Mesh is shown in *Image 1*. It is seen from Top-View (**Num-7**). Now we're going to insert five Shapes.

Select the mesh in Object mode.

First create the "Basis" shape by pressing **i**->*Mesh*. (Or by pressing the Add Shape Key button under the Shapes tab in the Edit/Buttons window).

Note
Make sure that you have created/edited the mesh to your liking **before** you create the "Basis" shape. If you edit the mesh after creating shape-keys unpredictable results may occur.

We create the following four shapes directly at the beginning, so they all originate from the "Basis" shape. Repeat the **i**->*Mesh* process four times.

You can switch between the shapes by selecting the respective shape in the *Shapes* panel in the *Editing* buttons (**F9**). Now we're going to change the shapes.

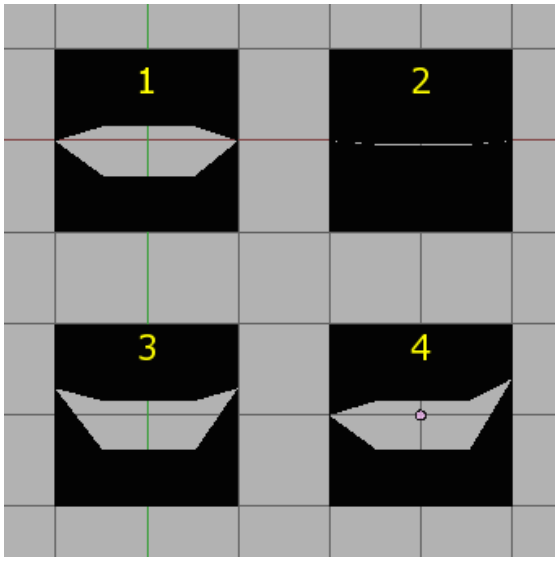


Image 2: Four of the five different shapes: Basis (1), Close (2), Smile (3), LeftUp (4).
The shapes:

1. Basis - Basis position, Mouth half opened.
2. Close - Mouth closed.
3. Smile - Corners of mouth pointing upwards.
4. LeftUp - Only left corner upwards.
5. RightUp - Only right corner upwards.

See the images ---------->>>>>>>>>

To drive the Ipo curves we use the bones of an armature. One bone to open and close the mouth, another one to lift or lower the corners of the mouth. Since we can use all three directions in space separately as drivers input, we can move both corners of the mouth by moving in vertical direction, by moving in horizontal direction we move the right resp. left side only, etc.

It's up to you how many bones you use and how you want to control the shapes, it depends on the amount of control and ease of use you would like to achieve.

In our example the armature is inserted in ZX (Front) view (**Num-1**). You may insert it on it's own layer and use a separate 3D window to move the bones around. You should align the bones along the global axles, it is easier than to confine their movement.

Name the left bone "Close", the right bone "Smile". Name the armature object "Driver" (**F9**->*Link and Materials* panel, *OB:* field.). (The name "Driver" is arbitrary and has no relation to the IPO->"Driver")

Select the mesh in object mode.

Split the 3D window and change to the *Ipo Curve editor*. Select the *Ipo Type->Shape* (*Image 3*).

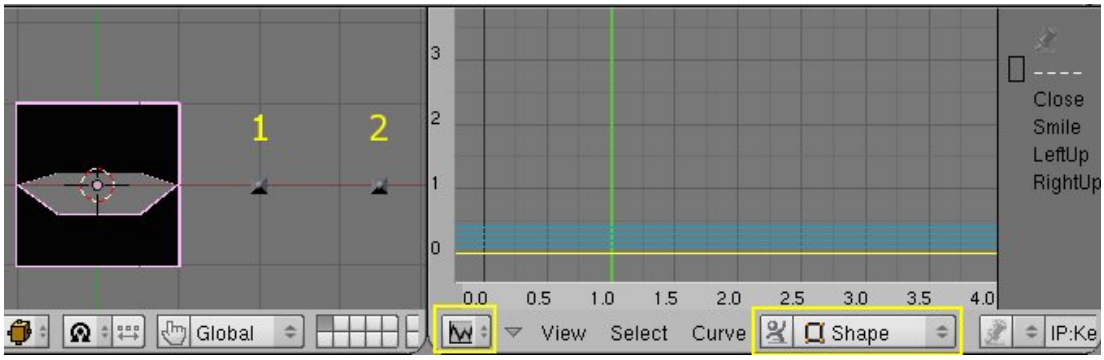


Image 3: Setup of the Armature. On the right side the Ipo window in Shape mode. We're looking at the bones bottom-up. 1: "CLOSE", 2: "Smile".

Connecting Shape key and Driver

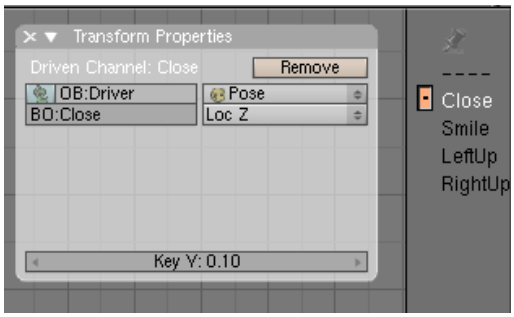


Image 4: A Bone as Ipo key Driver in the Ipo Curve editor.
To connect the Shape key with the controlling object - the Driver - you select the Shape key (LMB from the list on the right side of the IPO window), and press **n** in the Ipo window in order to call the *Transform Properties* panel. Click on *Add Driver*. Insert the name of the armature in the *OB:* field and select *Pose* in the drop down box. Insert the name of the controlling bone in the field *BO:* (here "Close"). We use *LocZ* to control the pose (*Image 4*). If you're in doubt about the correct coordinate, select the bone in the 3D window and use it's transform properties panel in the 3D window to watch it's coordinates.

To insert an Ipo curve press **i** in the Ipo Curve editor and confirm *Default one-to-one mapping*. A linear curve transition from (0,0) to (1,1) is inserted. The symbol for the Ipo curve in front of the shapes name receives a point, in order to indicate that a Driver for the curve is present. The Ipo curve maps the movement of the bone to the *Key value* of the shape, e.g. if you move the bone one Blender Unit (BU) in Z direction, the *Key value* of the shape changes from 0 to 1. If you move in negative Z direction the shape is inverted, i.e. the vertices move in opposite direction. If the Ipo curve runs horizontally the shape does not change if you move the bone.

We will use that to specify the borders of a meaningful movement of the bone.

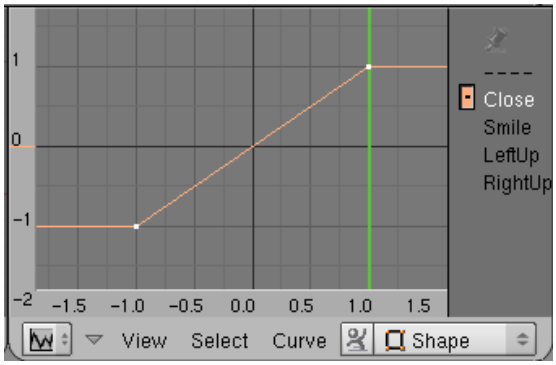


Image 5: Ipo curve for the shape "Close".
Set the mapping to (-1, -1), (1, 1) by selecting the curve and pressing â Tab to enter edit mode. Right click on the point at (0,0), press G to grab it, and move your mouse down and to the left until the point has shifted to (-1,-1). You also can, once the point is selected, directly type its new coordinates in the fields Vertex X: and Vertex Y:, at the bottom of the Transform Properties pannel. Press Tab again to leave edit mode. Set the Extend mode to *Constant* (*Curve->Extend mode->Constant*). Press **T**->*Linear* reset to a linear curve, if the type of the curve changes when the extend mode changes.

Repeat these steps for the shape "Smile", but we have to invert the Ipo curve (**S-X**) from (-1,1) to (1,-1), so that the movement of the bone corresponds with the movement of the mouth. In order to raise the corners of the mouth individually, the Shape key is connected with the x-coordinate of the Bone. If the Bone moves in positive x-direction the right corner is to be raised, if it moves in negative y-direction the left corner (*Image 6*).

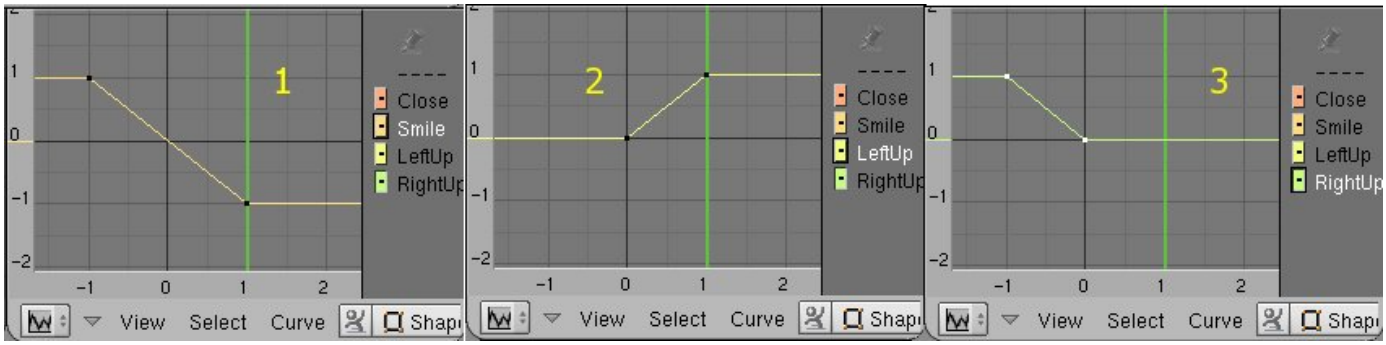


Image 6: The three Ipo curves for the Shapes "Smile" (1), "LeftUp" (2) and "RightUp" (3).
If you now move the bones in pose mode, the shapes change accordingly. If you need more shapes, or want the handling more flexible, you have to define more drivers.

Custom Controls

"Custom controls" - thus essentially selfmade manual controllers - facilitate the handling of Driven Shape Keys. There is no such Blenderobject (as of now), so we use mesh objects to visualise the purpose and borders of driver objects. The movement of the drivers is constricted by setting the lock options in their *Transform Properties* panel and - if you like - by floor constraints. So the driver objects behave pretty much like real control buttons.

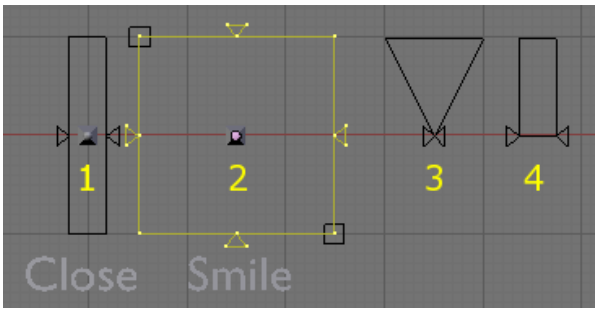


Image 7a: Four meshes to show the borders of the settings.

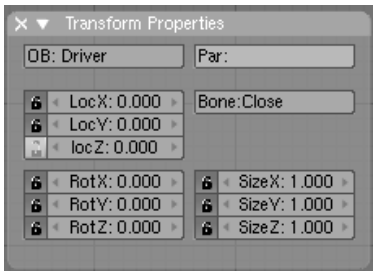


Image 8: Transform Properties panel for the bone "Close" (No. 1 in Image 7a.) The bone shall only move up and down.

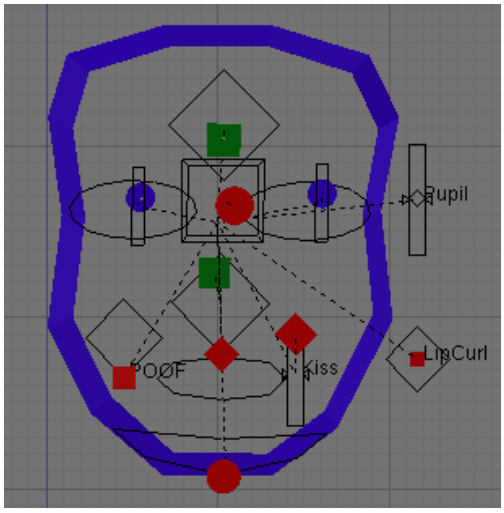


Image 7b: A more intuitive way to visualise the driver movements.

Image 7a shows some examples of Custom Controls (based on an idea by "PolygoneUK" from the Elysiun forum). By restricting the movement directions in the *Transform Properties* panel we allow the movement of the bone up and down, but not sideways.

You can go for a different visualisation if you build a stylized "Face", and insert the drivers in their respective places (idea by: Zeyne 2.4 rig with facial controller).

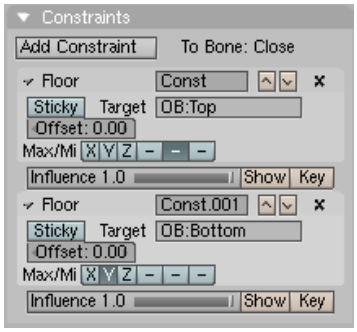


Image 9: If you use Floor Constraints you can confine the movements of the drivers to the area of the control. In example 2 two further objects at the corners are used as *Floor* constraints. They prevent the movement of the bones beyond the indicated borders. However, one needs four constraints for an area, only two constraints for the "Up and Down" sliders (*Image 9*).

After creating the controls we can move them to a global scene, so you don't edit them inadvertently. Add a new, empty scene (the double arrow besides the *SCE:* field, *Add new->Empty*). Name this new scene "Global". Mark all objects which belong to the control, but not the Armature, which we still have to change. Link the objects into the scene "Global" with CtrlL->*To scene... ->Global*.

Then we remove the connection between the objects: **U**->*Object*. Now delete the objects in the current scene. Change into the *Scene* Buttons (**F10**) and click on the small Button with the double arrow and the inconspicuous Tooltip *Scene to link as a Set* in the *Output* panel. Select "global". Now the controls are indicated, but you can't select or edit them.

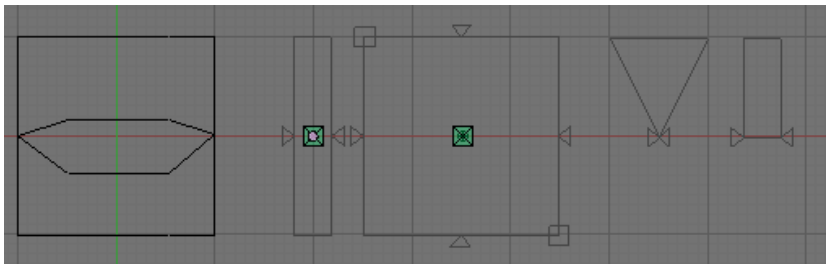


Image 10: The controls have been moved to a global scene.

NLA

Now the individual "face" movements can be arranged to actions, and these again in the NLA editor to complex groups. An action would be a certain face expression, e.g. merrily, sadly, surprised, bad etc. If one has provided all necessary actions, one works only with these actions.

Of course your're not forced to work in the NLA editor, instead you could create one long action, into which you insert all Ipos.

Define Actions

As described in the section The Non Linear Animation Editor Window, the different actions for the face have to be created.



Image 11: Three actions in the Action Editor window.
The example "face" is extended by two "eyes". Three actions were created: "Wink", "Joy" and "Anger" (*Image 11*). By positioning the bones of the armature in pose mode the expressions are created, all necessary bones selected and an Ipo inserted. Move ten frames foreward, and save exactly the same pose. The timing is made in the *NLA Editor* window.

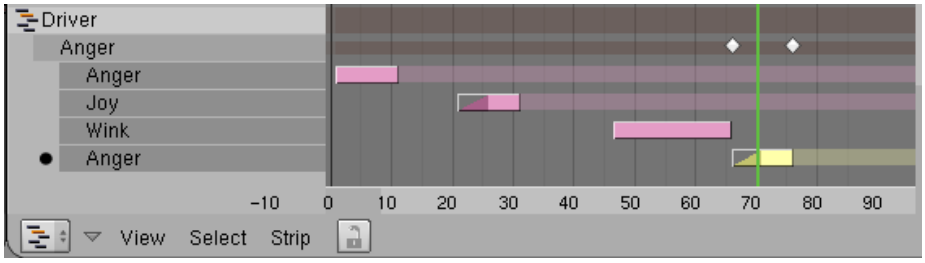


Image 12: The actions in the NLA Editor window.
Now arrange the expressions in the NLA editor and define the timing. Especially useful is the blending of different expressions (*Blendin/Blendout*). If you use the option *Hold* for the strips, you should define an action "Basis", so that you can blend to the "Basis" shape.

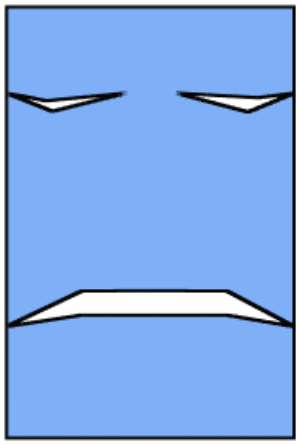


Image 13: Enjoy Blending!

Links

Ipo Drivers, Releasenotes of v2.40
Orange Blog: License to drive Driven Shape Keys, blend file provided.

Introduction

This chapter regroups various Blender advanced effects, often used to simulate real physical phenomena.

The Blender particle systems can be used to simulate many things: hair, grass, smoke, flocks, â←¦

Another great tool is the soft body simulator, useful for everything that tends to bend, deform, in reaction to forces like gravity or wind, or when colliding with other objectsâ←¦ It can be used for skin, rubber, and even clothes, even though there is a special tool for the latter case.

Blender also ships with some more general physics tools, like force fields, which might impact on particles, soft and rigid bodies, etc.

And you have a way to simulate very realistic fluids, with flows, obstacles, and so on!

Have fun, and Happy Blending!

Particles are lots of items emitted from mesh objects, typically in the thousands. Each particle can be a point of light or a mesh, and be joined or dynamic. They may react to many different influences and forces, and have the notion of a lifespan. Dynamic particles can represent fire, smoke, mist, and other things such as dust or magic spells. Static particles form strands and can represent hair, grass and bristles. You see particles as a Particle modifier, but all settings are done in the Particle sub-context of the Object context (F7).

Particles do not work with an Array modifier or with Group objects, it is planed to change that with version 2.5.

Incompatibility with Prior Versions
There are many differences between the â← oldâ← particle system that was used up to and including version 2.45, and the â← newâ← particle system. There are many things possible now that could not be done with the old system. The new system is incompatible to the old system, though Blender tries to convert old particle systems, which works only to some extent. The old system is most like the new Emitter system (keep reading to find out what that is). If you are using an old version of Blender 2.45 and previous, click here to access the old documentation.

Description

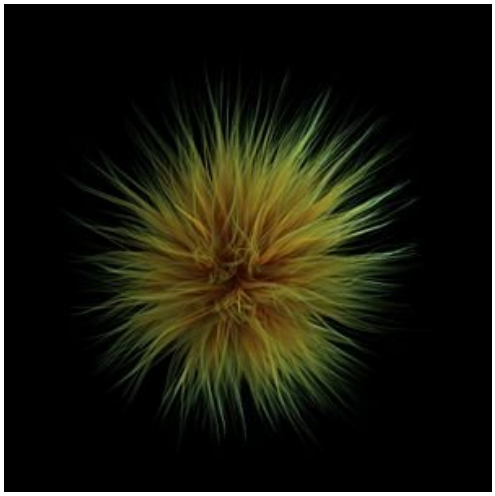


Image 1: Some fur made from particles (Blend file).
Particles generally flow out from their mesh into space. Their movement can be affected by many things, including:

- Initial velocity out from the mesh.
- Movement of the emitter (vertex, face or object) itself.
- Movement according to â← gravityâ← or â← air resistanceâ← .
- Influence of force fields like wind, vortexes or guided along a curve.
- Interaction with other objects like collisions.
- Partially intelligent members of a flock (herd, school, â←¦), that react to other members of their flock, while trying to reach a target or avoid predators.
- Smooth motion with softbody physics (only Hair particle systems).
- Or even manual transformation with Lattices.

Particles may be rendered as:

- Halos (for Flames, Smoke, Clouds).
- Meshes which in turn may be animated (e.g. fish, bees, â←¦). In these cases, each particle â← carriesâ← another object.
- Strands (for Hair, Fur, Grass); the complete way of a particle will be shown as a strand. These strands can be manipulated in the 3D window (combing, adding, cutting, moving, etc).

Every object may carry many particle systems. Each particle system may contain up to 100.000 particles. Certain particle types (Hair and Keyed) may have up to 10.000 children for each particle (children move and emit more or less like their respective parents). The size of your memory and your patience are your practical boundaries.

You can have multiple sets of particle systems attached to an object, so, for example, you can have long fur as one system, and then very short fur mixed in. These particle systems can be mixed and matched and shared among many objects in your blend file.

The â← newâ← system is much more powerful than the old, but there is one thing you canâ← t do with the â← newâ← system: you may no longer create branched static particle systems e.g. for very simple shrubbery.

Workflow

The process for working with particle is:

1. Create the base mesh which will emit the particles. This mesh is not rendered by default, but the base material for the mesh is used to color the particles. Since a mesh can carry multiple materials, each particle system may have its own material.
2. Create one or more Particle Systems to emit from the mesh. Many times, multiple particle systems interact or merge with each other to achieve the overall desired effect.
3. Tailor each Particle Systemâ← s settings to achieve the desired effect.
4. Animate the base mesh and other particle meshes involved in the scene.
5. Define and shape the path and flow of the particles.
6. (Only applies to Hair particle systems): Sculpt the emitterâ← s flow (cut the hair to length and comb it for example).
7. Make final render and do physics simulation(s), and tweak as needed.

Creating a Particle System

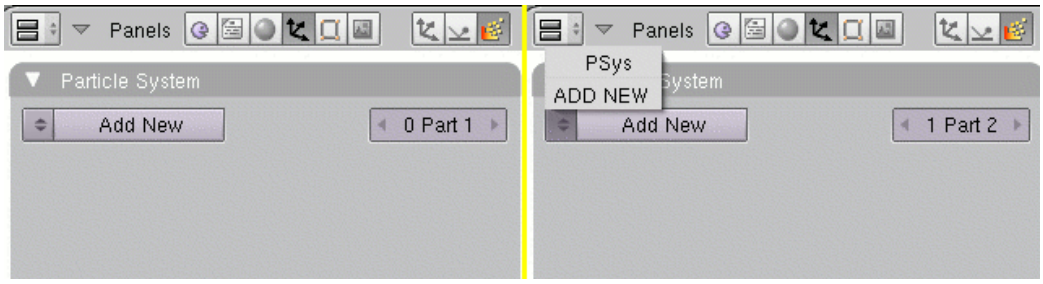


Image 2: Adding a particle system.
To add a new particle system to an object, go to the Particle sub-context of the Object (F7) context and click Add New in the Particle System tab. An object can have many Particle Systems. Therefore, when adding a Particle system, you may also:

- Select an existing particle system from the drop-down menu next to the Add New button.
- Advance the active particle system (X Part Y) selector to an empty slot by clicking it near its right edge. As with

material indices, the number to the left of Part indicates the active particle system, while the number on the right indicates the number of particle systems available for the object.

- Switch between already assigned Particle systems (X Part Y).
- Exchange one particle system for an other (for example `PA:PSys.002` to `PA:PSys.001`)

Types of Particle systems

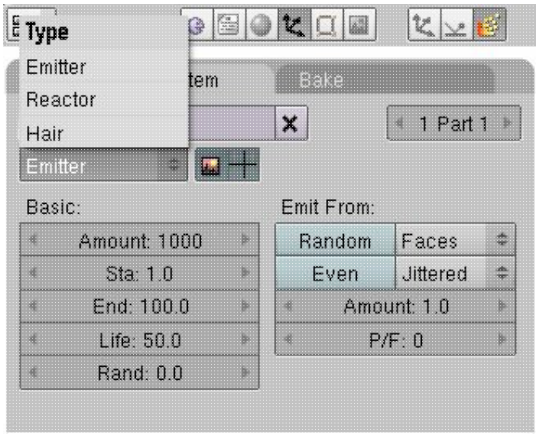


Image 3: Particle system types.
After you have created a particle system, the Buttons window fills with many panels and buttons. But donâ← t panic! There are three different types of particle systems, and you can change between these three with the Type drop-down list:

Emitter
: This parallels the old system to the greatest extent. In such a system, particles are emitted from the selected object from the Start frame to the End frame and have a certain lifespan. You may also render these kind of particles as strands (depending on the particleâ← s physics).

Reactor
: Reactor particles are born when other systemsâ← particles do things. Usually, but not with emit from particles, the target particleâ← s size determines itâ← s area of influence. This is useful for most of the effects where you would have used children in the old particle system.

Hair
: This system type is rendered as strands and has some very special properties: it may be edited in the 3D window in realtime and you can also animate the strands as Softbodies (similar to Softbody curves).

The settings in the Particle System panel are different for each system type. For example, in *Image 3* they are shown for only system type Emitter.

Common Options

Each system has the same basic sets of controls, but options within those sets vary based on the system employed. These sets of controls are:

- Particle System - basic controls about quantity and quality.
- Physics - how the particles behave.
- Visualization - how to see the particles, in the 3D viewport and when rendered.
- Extras - controlling emission, interaction and time.
- Children - particles spawning more particles.

Some individual settings are common for more than one particle system:

- Datablock selector (PA:): here you can select particle systems (browser button), name them (text field), or delete them (X button).
- X Part Y: selector for the active particle system.
- Enable buttons: toggle for enabling or disabling the active particle system in the 3D window or the rendering.
- Type: menu for selecting the type of particle system.

Bake

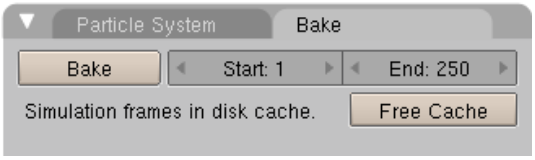


Image 4: Bake panel for particles.
Emitter and Reactor Systems use a unified system for caching and baking (together with softbody and cloth). The results of the simulation are automatically cached to disk when the animation is played, so that the next time it runs, it can play again quickly by reading in the results from the disk. If you Bake the simulation the cache is protected and you will be asked when youâ← re trying to change a setting that will make a recalculating necessary.

Beware of the Start and End Settings:

The simulation is only calculated for the positive frames in-between the Start and End frames of the Bake panel, whether you bake or not. So if you want a simulation longer than 250 frames you have to change the End frame!

- Caching
 - As animation is played, each physics system writes each frame to disk, between the simulation start and end frames. These files are stored in folders with prefix â← `blendcache`â← , next to the .blend file. Note that for the cache to fill up, one has to start playback before or on the frame that the simulation starts.
 - The cache is cleared automatically on changes - but not on all changes, so it may be necessary to free it manually e.g. if you change a force field.
 - If it is impossible to write in the subdirectory there will be no caching.
 - The cache can be freed per physics system with a button in the panels, or with the CtrlB shortcut key to free it for all selected objects.
 - If the file path to the cache is longer than what is possible with your operating system (more than 250 characters for example), strange things might happen.

- Baking
 - The system is protected against changes after baking.
 - The Bake result is cleared also with CtrlB for all selected objects or click on Free Bake for a singular particle system.
 - If the mesh changes the simulation is not calculated anew.
 - Sorry: no bake editing for particles like for softbodies and clothes.

Interaction in real time

To work with particle systems you may find it handy to use the Timeline window. You can change between frames and the particle system will always be shown in the actual state. The option Continue Physics in the Playback menu of the Timeline window lets you interact in real time with the particle system, e.g. by moving collision objects or shake an emitter object. And this is real fun!



Continue Physics does not work while playing the animation with AltA:

Right. This works only if you start the animation with the Play button of the Timeline window.

Two notes at the end:

- For renderfarms, it is best to bake all the physics systems, and then copy the blendcache to the renderfarm as well.
- Be careful with the sequence of modifiers in the modifier stack (as always). You may have a different number of faces in the 3D window and for rendering (from v2.47 upwards), if so, the rendered result may be very different from what you see in the 3D window.

We will continue with an overview over the three different particle system types.

Links

- Tutorials
- Physics Caching and Baking

- Particles Change Log
- Particle Rewrite Documentation
- Thoughts about the particle rewrite code
- Static Particle Fur Library

Authorâ← s Note
I have written the text and created the images for the german manual [1] --Soylentgreen 19:53, 4 February 2008 (CET)
I've taken quite a lot of text from BlenderDev/Particles_Rewrite_Doc and from [2]. So who is to be credited for that and in what form? --Soylentgreen 20:13, 5 February 2008 (CET)

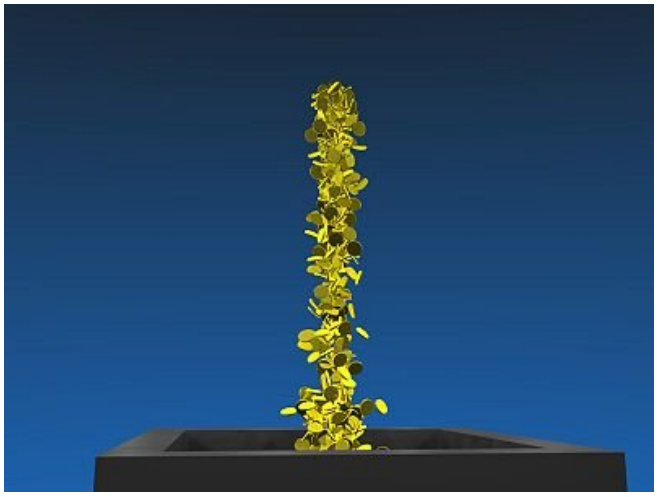


Image 1: A simple emitter particle system. Blend file
Animation of *Image 1* (info)

Emitter

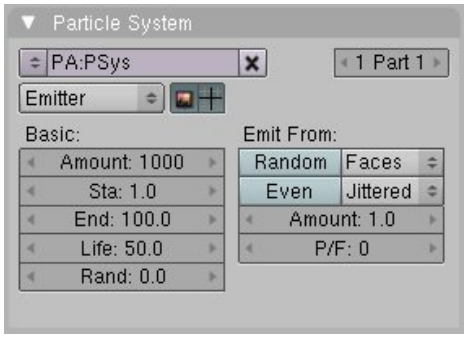


Image 2a: Settings for the Emitter particle system.
The Emitter system works just like its name says: it emits/produces particles for a certain amount of time. In such a system, particles are emitted from the selected object from the Start frame to the End frame and have a certain lifespan. These particles are rendered default as Halos, but you may also render these kind of particles as objects or strands (depending on the particleâ← s physics, see Visualization).

Options

Basic

- The two buttons next to the field type activate rendering resp. show the particles in the 3D window.
- Amount: The maximum amount of parent particles used in the simulation.
- Sta: The start frame of particle emission. You may set negative values from version 2.48 on. This enables you to start the simulation before the actual rendering.
- End: The end frame of particle emission. Remember that you have to set the Bake values also if you need the simulation in other frames than 1 to 250.
- Life: The lifetime (in frames) of the particles.
- Rand: A random variation of the lifetime of a given particle. The shortest possible lifetime is *Life*Ã (1-*Rand*). Values above 1.0 are not allowed. For example with the default Life value of 50 a Rand setting of 0.5 will give you particles with lives ranging from 50 frames to 50Ã (1.0-0.5)=25 frames, and with a Rand setting of 0.75 youâ← ll get particles with lives ranging from 50 frames to 50Ã (1.0-0.75)=12.5 frames.

If you want to control the emission over time, you can use an Ipo Curve (see Controlling Emission, Interaction and Time).

Emit From

These parameters define how the particles are emitted, giving precise control over their distribution. You may use vertex groups to confine the emission, that is done in the Extras panel. You may also control the particle emission with an (animated) texture.

- Emitter element: Verts/Faces/Volume: States that particles are emitted respectively by vertices, faces, or the enclosed volume of the mesh emitter.
- Random: The emitter element indices are gone through in a random order instead of linearly (one after the other).
- Even: Particle distribution is made even based on surface area of the elements, i.e. small elements emit less particles than large elements, so that the particle density is even.
- Distribution:
 - ♦ Jittered: Particles are placed at jittered intervals on the emitter elements.
 - ◊ Amount: Amount of jitter applied to the sampling.
 - ◊ P/F: Number of emissions per face (0 = automatic).
 - ♦ Random: Particles are placed randomly in the emitterâ← s elements.
 - ♦ Grid: Particles are set in a 3d grid and particles near/in the elements are kept.
 - ◊ Resol: Resolution of the grid.
 - ◊ Invert: Toggles what is considered to be the emitter.

Usage

You use an emitter system when you need a lot of identical moving elements, for example: a particle fire (Tutorial) where you use the particles to create flames and sparks. Some other applications where you would use an emitter are: smoke rising from a fire or cigarette, ants issuing from an anthill, bats flying out of a cave, etc.

Reactor

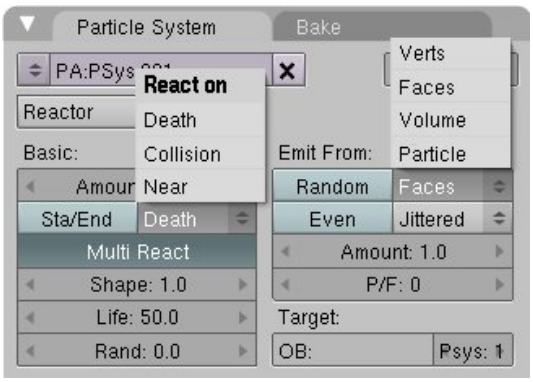


Image 3a: Settings for a Reactor particle system.
Using a Reactor implies you are using at least two particle systems, because reactor particles are born, or created, as a result of the actions of another particle system. The other particles may have come from another particle system on this object, or a particle system on a separate object.

The system that reacts is called the â← *Target*â← .

You should use a Reactor particle system in the places where you used children with the older implementation of particle physics.

By having the reactor produce particles upon the death of particles in the other system, you can create particle cascades.

Usually (not with emit from particles), the target particleâ← s size determines itâ← s area of influence (Size button in the Extras panel).

You should set up the particle systems starting with the last one in a chain, so if you have a cube with `pSysOriginator` that targets a sphere with `pSysTarget`, you should be sure to set up `pSysTarget` before you set up `pSysOriginator`. Systems that create a loop, for instance `pSysOriginator` â `pSysTarget` â `pSysOriginator`, will probably cause problems, since one target system wonâ← t be updated before the system that targets it is.

Options

Basic

- Sta/End: Particles are only emitted when the chosen event happens (see React on, below). But with this option, you can force all remaining particles to be emitted within the Sta through End frames (the proper settings are displayed when this option is active).
- React on: Which event of the target particles triggers emission: Death (target particle dies), Collision (target particle collides), or Near (target particle in the vicinity of the Reactor).
- Multi React: Allows reacting multiple times (alive particles react too and not just unborn particles).
- Shape: How reaction strength varies with distance from target particle. The closest, the strongest the reaction.

Emit From

These parameters are mostly the same as the Emitter particle system type, to one exception:

- Emitter element: Particle is another possible option for Reactor particle systems, meaning that particles are emitted by other particles, when reacted on.

Target

These parameters are only useful for Reactor systems, and are meant to define the particle system to which the Reactor system should react.

- OB: Defines the object that has the target particle system, whose particles are evaluated in search for events to react to. If the field is blank, the current object is used.
- Psys: Selects which particle system is the target. Should appear red when no valid target is specified. This is always the case when the current particle system is the first particle system of the current object (OB: field empty).

Usage

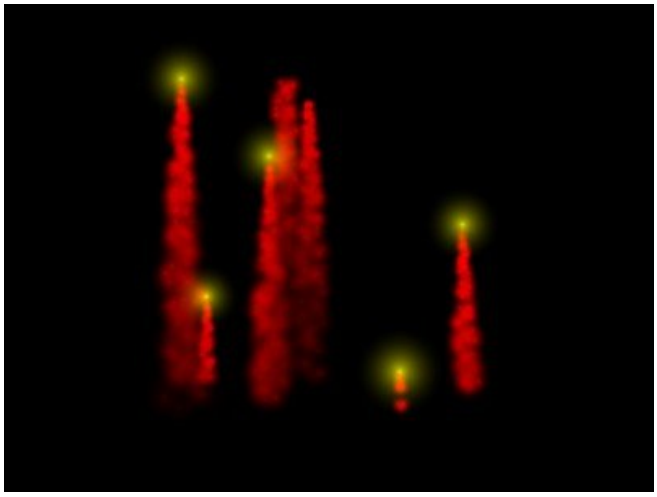


Image 3b: Example for reactor particles (red) reacting to near particles (yellow).
Animation of *Image 3b* (info)
Interesting examples of reactions:

- React on - Death: Particles are emitted when the target particles die. Two example usages:
 - ♦ Fireworks: Target particles are emitted upwards with normal gravity. Reactor particles are set to emit from particles with random initial velocity and similar gravity.
 - ♦ Minefield: Reactor particles emitted from a large volume with reactor initial velocity. Target particles that die inside the volume cause â← explosionsâ← when they die.
- React on - Collision: Particles are emitted when target particles collide with something. Example usage:
 - ♦ Raindrops: Reactor particles are on the ground plane with normal and reactor velocity. Target particles fall from above and collide with the ground plane.

- React on - Near: Particles are emitted when target particles are near them. Example usages:
 - ♦ Trails: Reactor particles set to emit from particles with random initial velocity. Now the reactor particles are always â← nearâ← the target particles so they emit constantly leaving a trail for the target particles (*Image 3b*).
 - ♦ Sand dunes: Reactor particles are emitted from the ground mesh with negative reactor velocity. When target particles fly over the ground they make reactor particles rise from the ground.

- Tutorial no. 1: Fireworks
- Tutorial no. 2: Fireworks.

Hair

This particle system creates only static particles, which may be used for hair, fur, grass and the like. Only this system may be edited interactively in the 3D window (in Particle Mode), only this system may be animated as softbody.

The complete path of the particles is calculated in advance. So everything a particle does a hair may do also. A hair is as long as the particle path would be for a particle with a lifetime of 100 frames. Instead of rendering every frame of the particle animation point by point there are calculated control points with an interpolation, the segments.

Options

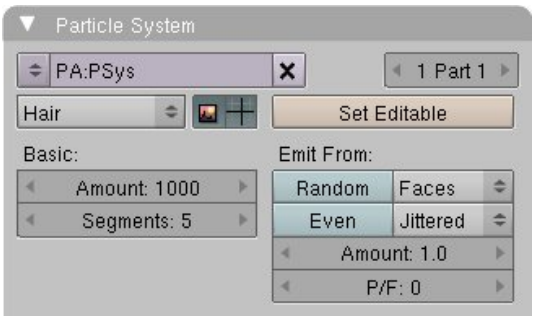


Image 4a: Settings for a Hair particle system.

- Set Editable: The system will become editable in Particle Mode. You canâ← t change the number of particles or the particle physics if you have set the hair editable. If you need to change these things later all changes in Particle Mode will be lost.
- Amount: Use as little particles as possible, especially if you plan to use softbody animation later. But you need enough particles to have good control. For a â← normalâ← haircut I found some thousand (very roughly 2000) particles to give enough control. You may need a lot more particles if you plan to cover a body with fur. Volume will be produced later with Children.
- Segments: The number of segments (control points minus 1) of the hair strand. In between the control points the segments are interpolated. The number of control points is important:
 1. for the softobdy animation, because the control points are animated like vertices, so more control points mean longer calculation times.
 2. for the interactive editing, because you can only move the control points (but you may recalculate the number of control points in Particle Mode).

10 Segments should be sufficient even for very long hair, 5 Segments are enough for shorter hair, and 2 or 3 segments should be enough for short fur.

Usage



Image 4b: Particle systems may get hairyâ←¦

- We deal with the production of longer hair on the page Hair.
- Fur Tutorial, which produced (*Image 4b*). It deals especially with short hair.

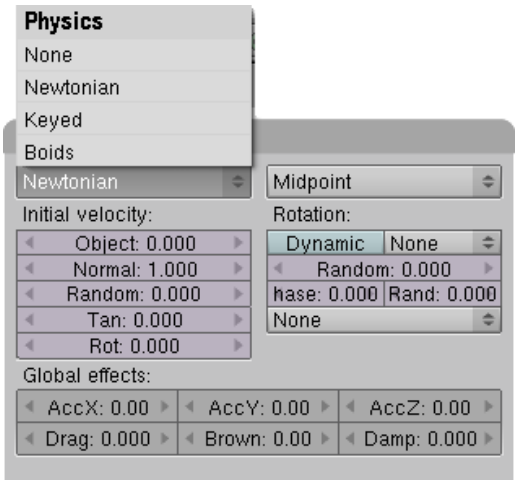


Image 1: The Physics panel for particles.
The movement of particles may be controlled in a multitude of ways:

- With particles physics: there are four different systems:
 - ♦ None: It doesnâ← t give the particles any motion, which makes them belong to no physics system.
 - ♦ Newtonian: Movement according to physical laws.
 - ♦ Keyed: Dynamic or static particles where the (animated) targets are other particle systems.
 - ♦ Boids: Particles with limited artificial intelligence, including behavior and rules

programming, ideal for flocks of birds or schools of fishes, or predators vs preys simulations.
- By softbody animation (only for Hair particle systems).
- By forcefields and along curves.
- By lattices.

Here we will discuss only the particle physics in the narrower sense, i.e. the settings in the Physics panel.

Physics: None

At first a Physics type that makes the particles do nothing could seem a bit strange, but it can be very useful at times. None physics make the particles stick to their emitter their whole life time. The initial velocities here are for example used to give a velocity to particles that are effected (or affected?) by a harmonic effector with this physics type when the effect of the effector ends.

Moreover, it can be very convenient to have particles at disposal (whose both Unborn and Died are visible on render) to groom vegetation and/or ecosystems using Object, Group or Billboard types of visualization.

Physics: Newtonian

These are the â← normalâ← particle physics. Particles start their life with the specified initial velocities and angular velocities, and move according to forces. The response to environment and to forces is computed differently, according to any given integrator chosen by the animator.

Integrators

Integrators are a set of mathematical methods available to calculate the movement of particles. The following guidelines will help to choose a proper integrator, according to the behavior aimed at by the animator.

- Euler: Also known as â← Forward Eulerâ← . Simplest integrator. Very fast but also with less exact results. If no dampening is used, particles get more and more energy over time. For example, bouncing particles will bounce higher and higher each time. Should not be confused with â← Backward Eulerâ← (not implemented) which has the opposite feature, energies decrease over time, even with no dampening. Use this integrator for short simulations or simulations with a lot of dampening where speedy calculations is more important than accuracy.
- Midpoint: Also known as â← 2nd order Runge-Kuttaâ← . Slower than Euler but much more stable. If the acceleration is constant (no drag for example), it is energy conservative. It should be noted that in example of the bouncing particles, the particles might bounce higher than they started once in a while, but this is not a trend. This integrator is a generally good integrator for use in most cases.
- RK4: Short for â← 4th order Runge-Kuttaâ← . Similar to Midpoint but slower and in most cases more accurate. It is energy conservative even if the acceleration is not constant. Only needed in complex simulations where Midpoint is found not to be accurate enough.

Initial velocity

The initial velocity of particles can be set through different parameters, based on the type of the particle system (see Particle System tab). If the particle system type is Emitter or Hair, then the following parameters give the particle an initial velocity in the direction ofâ←¦

- Object: â←¦the emitter objects movement (i.e. let the object give the particle a starting speed).
- Normal: â←¦the emitterâ← s surface normals (i.e. let the surface normal give the particle a starting speed).
- Random: â←¦a random vector (i.e. give the starting speed a random variation in direction and in value. You can use a texture to only change the value, see Controlling Emission, Interaction and Time).
- Tan & Rot: â←¦a tangential vector along the surface, rotated by Rot.
 - ♦ Tan: Let the tangent speed give the particle a starting speed.
 - ♦ Rot: Rotates the surface tangent.

If the particle system type is Reactor, then the following parameters give the particle an initial velocity in the direction ofâ←¦

- Particle: â←¦the target particles velocity (i.e. let the target particle give the particle a starting speed).
- Reactor: â←¦a vector away from the target particles location at the time of the reaction (i.e. let the vector away from the target particles location give the particle a starting speed).

Rotation

These parameters specify how the individual particles are rotated during their travel. To visualize the rotation of a particle you should choose visualization type Axis in the Visualization panel and increase the Draw Size.

- Dynamic: If enabled, only initializes particles to the wanted rotation and angular velocity and letâ← s physics handle the rest. Particles then change their angular velocity if they collide with other objects (like in the real world due to friction between the colliding surfaces). Otherwise the angular velocity is predetermined at all times (i.e. set rotation to dynamic/constant).
- Rotation: Sets the initial rotation of the particle by aligning the x-axis in the direction ofâ←¦
 - ♦ None: â←¦the global x-axis.
 - ♦ Normal: â←¦the emitterâ← s surface normal.
 - ♦ Velocity: â←¦the particleâ← s initial velocity.
 - ♦ Global X/Global Y/Global Z: â←¦one of the global axes.
 - ♦ Object X/Object Y/Object Z: â←¦one of the emitter object axes.
- Random: Randomizes rotation.
- Phase/Rand: Initial rotation phase, Rand allows a random variation of the Phase.
- Angular v: The magnitude of angular velocity, the dropdown specifies the axis of angular velocity to beâ←¦
 - ♦ None: â←¦a zero vector (no rotation).
 - ♦ Spin: â←¦the particles velocity vector.
 - ♦ Random: â←¦a random vector.

If you use a Curve Guide and wanâ← t the particles to follow the curve, you have to set Angular Velocity to Spin and leave the rotation on Constant (i.e. donâ← t turn on Dynamic). Curve Follow does not work for particles.

Global effects

These parameters specify global physical factors that accelerate or decelerate particles velocity. Useful when simulating various phenomenon like gravity, air-drag, friction and such. Other, more complicated or localized forces can be created with Force Fields.

- AccX, AccY and AccZ: An acceleration in the direction of the global axes. Use this to implement gravity by setting AccZ to a negative value, for example.
- Drag: A force that reduces particle velocity in relation to itâ← s speed and size (useful in order to simulate Air-Drag or Water-Drag).
- Brown: A random force that changes from frame to frame. Simulates Brownian movement which is an effect seen on (very) small particles where forces from individual molecules are unbalanced over time. This is nice to simulate small, random wind forces.
- Damp: Reduces particle velocity (deceleration, friction, dampening).

Physics: Keyed

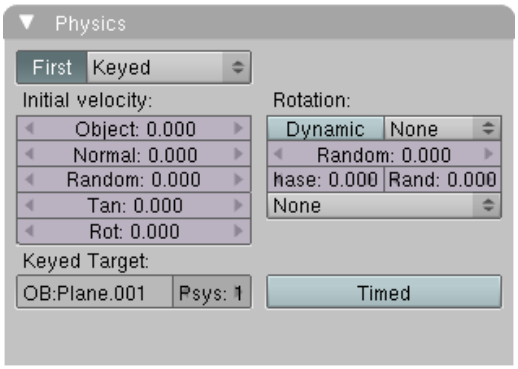


Image 2: The first of a chain of keyed particle systems.
The particle paths of keyed particles are determined from the emitter to another particle systemâ← s particles. This allows creation of chains of systems with keyed physics to create long strands or groovy moving particles. Basically the particles have no dynamics but are interpolated from one system to the next at drawtime.

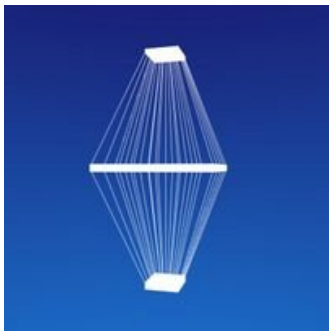


Image 3: Keyed Particles allow for great control and complex animations.
Animation for *Image 3* (info)
Because you have so much control over these kind of systems, you may use it e.g. for machines handling fibers (animation of a loom, â←¦). In (*Image 3*), the strands flow from the bottom system (*First keyed*) to the second keyed system in the middle, and from that to the top system that has None-Physics. Since you may animate each emitter object as you like, you can do arbitrarily complex animations.

To setup Keyed particles you need at least two particle systems.

- The first system has keyed physics, and it needs the option First activated. This will be the system thats is visible. *The second system may be another keyed system but without the option First, or a normal particle system. This second system is the target of the keyed system.

- Keyed Target: You have to enter the name of the object which bears the target system and if there are multiple particle systems the number of the system.

If you use only one keyed system the particles will travel in their lifetime from the emitter to the target. A shorter lifetime means faster movement. If you have more than one keyed system in a chain, the lifetime will be split equally. This may lead to varying particle speeds between the targets.

- Timed: This option is only available for the first keyed system. It works together with the Time slider for the other keyed systems in a chain.
- The Time slider allows to define a fraction of particle lifetime for particle movement.

An example: letâ← s assume that you have two keyed systems in a chain and a third system as target. The particle lifetime of the first system shall be 50 keys. The particles will travel in 25 frames from the first keyed system to the second, and in further 25 frames from the second system to the target. If you use the Timed button for the first system, the Time slider appears in the second systems panel. Itâ← s default value is 0.5, so the time is equally split between the systems. If you set Time to 1, the movement from the first system to the second will get all the lifetime (the particles will die at the second system).

If you set Time to 0 the particles will start at the second system and travel to the target.

Physics: Boids

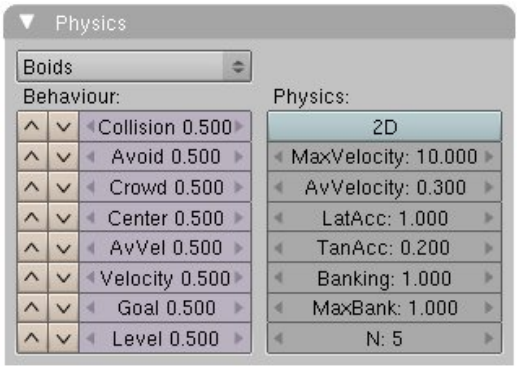


Image 4: Physics panel for Boids particle.
Boids particle systems can be set to follow basic rules and behaviors. They are useful for simulating flocks, swarms, herds and schools of various kind of animals, insects and fishes. They can react on the presence of other objects and on the members of their own system. Boids can handle only a certain amount of information, therefore the sequence of the Behaviour settings is very important. In certain situations only the first three parameter are evaluated.

Boids try to avoid objects with activated Deflection. They try to reach objects with positive Spherical fields, and fly from objects with negative Spherical fields. The objects have to share one common layer to have effect. It is not necessary to render this common layer, so you may use invisible influences.

Behaviour

- Only a certain amount of information can be evaluated. If the memory capacity is exceeded, the remaining rules are ignored.
- The rules are parsed from top-list to bottom-list (thus giving explicit priorities), and the exact order can be modified using the little arrows in front of each row. The list of rules available are:
 - ♦ Collision: Avoid objects with activated Deflection.

- Avoid: Avoid â← predatorsâ← (objects with Spherical fields and negative Strength).
- Crowd: Avoid other boids.
- Center: Get to flock center.
- AvVel: Maintain average velocity.
- Velocity: Match velocity of nearby boids.
- Goal: Seek goal (objects with Spherical fields and positive Strength).
- Level: Keep the Z level. The boids then try to not change their flightlevel. This is deactivated for 2D boids.

Each rule can be individually weighted; the value should be considered how hard the boid will try to respect a given rule (a value of 1.000 means the Boid will always stick to it, a value of 0.000 means it will never). If the boid meets more than one conflicting condition at the same time, it will try to fulfill all the rules according to the respective weight of each. Any rule could be weighted from -1.000 to +2.000 in order to give it more or less significance.

- Normal behavior can be expected with weights between 0.000 to 1.000.
- From 1.000 to 2.000 the boids over react according to the rules.
- From -1.000 to 0.000 the boids react contrary to the rules.

Please note that a given boid will try as much as it can to comply to each of the rules he is given, but it is more than likely that some rule will take precedence on other in some cases. For example, in order to avoid a predator, a boid could probably â← forgetâ← about Collision, Crowd and Center rules, meaning that â← while panickedâ← it could well run into obstacles, for example, even if instructed not to, most of the time.

As a final note, the Collision algorithm is still not perfect and in research progress, so you can expect wrong behaviors at some occasion. It is worked on.

Physics

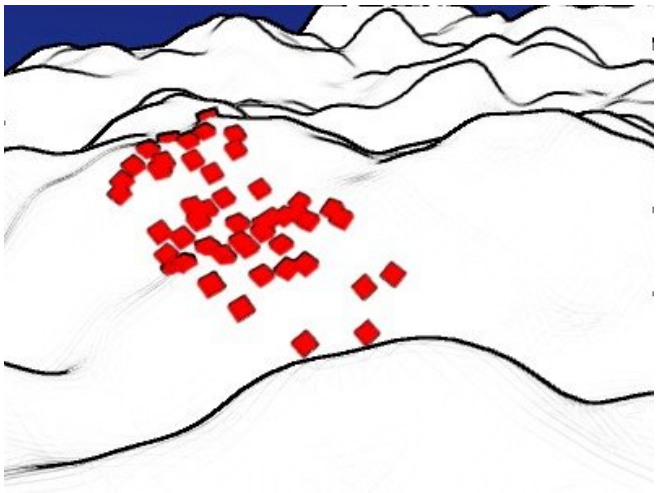


Image 5: Boids particle are able to follow a curved surface.

Animation for *Image 5* (info)

- MaxVelocity: Maximum velocity.
- AvVelocity: The usual speed percent of max velocity. If MaxVelocity is set to 10.000 and AvVelocity to 0.300, then the average velocity of the boids is 3.000.
- LatAcc: Lateral acceleration percent of max velocity (turn). Defines how fast a boid is able to change

direction.
- TanAcc: Tangential acceleration percent of max velocity (forward). Defines how much the boid can suddenly accelerate in order to fulfill a rule.
- Banking: Banking of boids on turns (1.0 == natural banking).
- MaxBank: How much a boid can bank at a single step.
- N: How many neighbours to consider for each boid.
- 2D: Constrains boid to a surface: either to the surface of a given object (if specified in the OB field) or to a certain Z value (GroundZ). Useful to simulate herds on a ground, for example. When activated, Level, Banking and MaxBank become irrelevant.
 - ♦ GroundZ: Default Z value.
 - ♦ OB: Objectâ← s surface the boid is constrained to.
 - ♦ If boids trajectory leads them out of the surface of an object, the GroundZ value is then used. E.g. Boids will distribute on the top half of a sphere and then â← dripâ← to the ground.

Boids, deflectors and effectors

As mentioned before, very much like Newtonian particles, Boids will react to the surrounding deflectors and fields, according to the needs of the animator:

- Deflection: Boids will try to avoid deflector objects according to the Collision ruleâ← s weight. It works best for convex surfaces (some work needed for concave surfaces).

For boid physics, Spherical fields define the way the objects having the field are seen by others. So a negative Spherical field (on an object or a particle system) will be a predator to all other boids particle systems, and a positive field will be a goal to all other boids particle systems.

When you select an object with a particle system set on, you have in the Fields tab a little menu stating if the field should apply to the emitter object or to the particle system. You have to select the particle system name if you want prey particles to flew away from predator particles.

- Spherical fields: These effectors could be predators (negative Strength) that boids try to avoid or targets (positive Strength) that boids try to reach according to the (respectively) Avoid and Goal rulesâ← weights. Sphericalâ← s effective Strength is multiplied by the actual relevant weight (e.g. if either Strength or Goal is null, then a flock of boids wonâ← t track a positive Spherical field).

You can also activate Die on hit (Extras panel) so that a prey particle simply disappears when â← attackedâ← by a predator particle which reaches it. To make this work, the predator particles have to have a spherical field with negative force, it is not sufficient just to set a positive goal for the prey particles (but you may set the predators force strength to -0.01). The size of the predators and the prey can be set with the Size button in the Extras panel.

Links

- Tutorial showing how to set a prey-predator relationship using Boids
- Boids in action
- Boids: Background and Update
- Flocks, Herds, and Schools: A Distributed Behavioral Model

With the items in the Visualization panel you can set the way the particles will be rendered or depicted in the view ports in various ways. Some option are valid only for the 3D window, the particles then are rendered always as Halos. Some of the options will be rendered as shown in the 3D window.



The emitter is invisible:

If you create a particle system, the emitter is no longer rendered. Activate the button Emitter to also render the mesh.

It is very often necessary to animate the material settings of the particle material. By default 100 frames of the Ipo curves are related to the lifetime of the particles. So a material animation of 100 frames in the Ipo window will take place in the lifetime of the particles, independent whether the particle life will last 10 or 1000 frames. You may change this relation in the Extras panel.

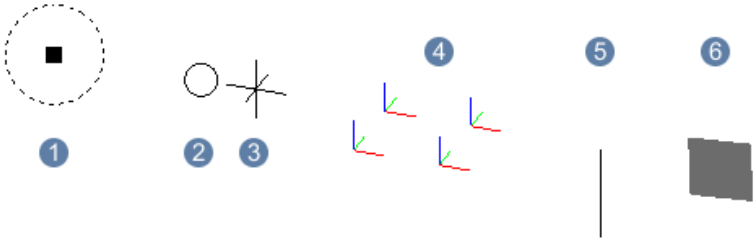


Image 1: Visualization types for particles.
In the 3D window particles can be depicted as:

1. Point, here shown with itâ← s size
2. Circle
3. Cross
4. Axis
5. Line or Path
6. Billboard

Object and Group visualizations are not shown.

Types of Visualization

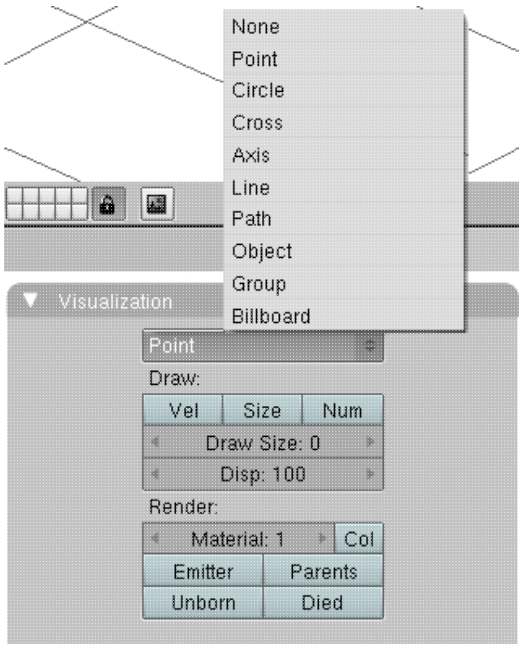


Image 2: The Visualization panel for particles.

- None: The particles are not shown in the 3D window and are not rendered. The emitter may be rendered though.
- Point/Circle/Cross: Particles visualized like Point, Circle, Cross and Axis are all still rendered as Halos, but displayed as needed in the 3D window. These modes of visualization donâ← t have any special options, but can be very useful when you have multiple particle systems at play, if you donâ← t want to confuse particles of one system from another (e.g. in simulations using Boids physics).
- Axis is useful if you want to see the orientation and rotation of particles in the view port. Increase the Draw Size until you can clearly distinguish the axis.
- Line: The Line visualization mode creates (more or less thin) polygon lines with the strand renderer in the direction of particles velocities. The thickness of the line is set with the parameter Start of the Strands shader (Material sub-context, Links and Pipeline panel).
 - ♦ Speed: Multiply the line length by particlesâ← speed. The faster, the longer the line.
 - ♦ Back: Set the length of the particleâ← s tail.
 - ♦ Front: Set the length of the particleâ← s head.
- Path: The way of the particle during itâ← s lifetime is shown at once. This is the visualization type you use for hair, grass etc. Needs either a Hair particle system or Keyed particle physics. Because this visualization type has so much options it is explained in greater detail below.
- Object: In the Object visualization mode the specified object (OB: field) is duplicated in place of each particle. The duplicated object has to be at the center of the coordinate system, or it will get an offset to the particle.
 - ♦ OB: The name of the object.
- Group: In the Group visualization mode, the objects that belong to the group (GR: field) are duplicated sequentially in the place of the particles.
 - ♦ GR: The name of the group.
 - ♦ Dupli Group: Use the whole group at once, instead of one of its elements, the group being displayed in place of each particle.
 - ♦ Pick Random: The objects in the group are selected in a random order, and only one object is displayed in place of a particle.

 Please note that this mechanism fully replaces old Blender particles system using parentage and DupliVerts to replace particles with actual geometry. This method is fully deprecated and doesnâ← t work anymore.

- Billboard: Billboards are aligned square planes. How they are aligned, and what are they aligned to can be influenced in many ways. Texturing billboards (including animated textures with alpha) is done by using uv coordinates that are generated automatically for them. This works well for animations, because the alignment of the billboards can be dynamic. An interesting alternative to billboards are in certain cases strands, because you can animate the shape of the strands. Because this visualization type has so much options it is explained in greater detail below.

Draw

Defines various display options for the particles in the viewports of Blender.

- Vel: Draw the velocity of the particles with a line.
- Size: Draw the size of the particles with a circle.
- Num: Draw the id-numbers of the particles in the order of emission.
- Draw size: Specifies how large (in pixels) the particles are drawn in the viewport (0 = default).
- Disp: Specifies the percentage of all particles to show in the viewport (all particles are still rendered).

Render

Defines various options for the particles mostly at render time.

- Material: Which objectâ← s material index to use for the particles.
- Col: Draw particles in the materialâ← s diffuse color (Col in the Material panel of the Material sub-context). This makes it easy to distinguish different particle systems in the 3D window. Please note the active particles system is always drawn in white in the viewports.
- Emitter: Render also the emitter object and not just its particles.
- Parents: Render also parent particles if child particles are used. Children have a lot of different deformation options, so the straight parents would stand between their curly children. So by default Parents are not rendered if you activate Children.
- Unborn: Render particles even before they are born.
- Died: Render particles even after they have died. This is very useful if particles die in a collision (Die on hit), so you can cover objects with particles.

Path

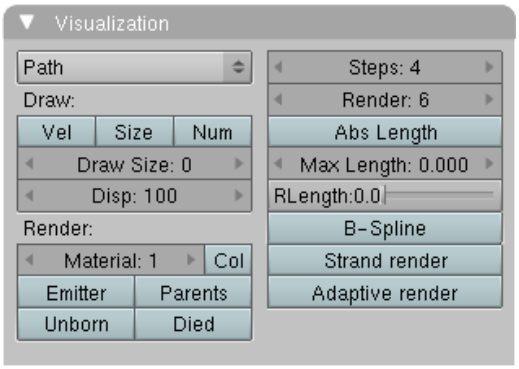


Image 3: The Visualization panel for Path visualization.
The Path visualization needs a Hair particle system or Keyed particles. It uses the strand renderer and is used for hair, fur, etc.

- Steps: Set the number of subdivisions of the drawn paths in the viewport (the value is a power of 2). This means 0 steps give 1 subdivision, 1 give 2 subdivisions, 2â 4, 3â 8, 4â 16, â←¦ *n*â 2 *n*.
- Render: Set the number of subdivisions of the rendered paths (the value is a power of 2). You should set this value carefully, because if you increase the render value by two you need four times more memory to render. Also the rendering is faster if you use low render values (sometimes drastically). But how low you can go with this value depends on the waviness of the hair.
- Abs Length: Use an absolute length for particle children path visualization mode.
 - ♦ Max Length: Set the value of the absolute maximum length (in blender units).
- RLength: Randomize path lengths.
- B-Spline: Interpolate hair using B-Splines. This may be an option for you if you want to use low Render values. You loose a bit of control but gain smoother paths.
- Strand render: [Keypointstrands] Use the strand primitive for rendering. Very fast and effective renderer.
 - ♦ Angle: How many degrees (0-45Â°) path has to curve to produce another render segment.
- Adaptive render: Tries to remove unnecessary geometry from the paths before rendering particle strands in order to make the render faster and easier on memory.
 - ♦ Angle: How many degrees path has to curve to produce another render segment (straight parts of paths need fewer segments).
 - ♦ Pixel: How many pixels path has to cover to produce another render segment (very short hair or long hair viewed from far away need fewer parts).

Please see also the manual page about Strands for an in depth description.

Billboard

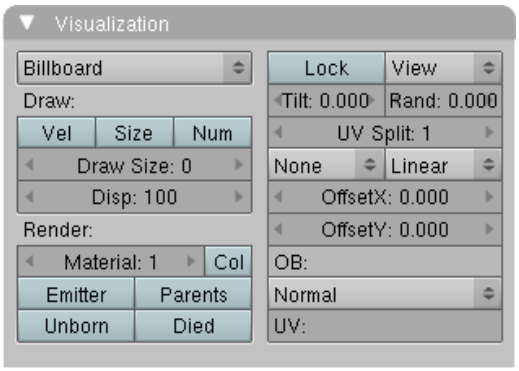


Image 4: Billboard visualization for particles.
Billboards are aligned square planes. They are aligned to the camera by default, but you can choose another object that they should be aligned to. If you move a billboard around itâ← s target, it always faces the center of itâ← s target. The size of a billboard is set with the parameter Size of the particle (in Blender Units). You can use them e.g. for Sprites, or to replace Halo visualization. Everything that can be done with a halo can also be done with a billboard. But billboards are real objects, they are seen by raytracing, they appear behind transparent objects, they may have an arbitrary form and receive light and shadows. They are a bit more difficult to set up and take more render time and resources.

Billboards can be textured (including transparency), so they can take an arbitrary shape. The textures can be animated in several ways:

- Depending on the particle lifetime (relative time).

- Depending on the particle starting time.
- Depending on the frame (absolute time).

You can use different sections of an image texture:

- Depending on the lifetime of the billboard.
- Depending on the emission time.
- Depending on align or tilt.

Since you use normal materials for the billboard you have all freedoms in mixing textures to your liking. The material itself is animated in absolute time.

The main thing to understand is that if the object doesnâ← t have any *UV Layers*, you need to create at least one in the objects Editing context, for any of these to work. Moreover, the texture has to be set to UV coordinates in the Map Input panel. If you want to see examples for some of the animation possibilities, see the Billboard Animation Tutorial.

Options

- Align to/Lock: You can limit the movement with these options. How the axis is prealigned at emission time.
 - ♦ View: No prealignement, normal orientation to the target.
 - ♦ X/Y/Z: Along the global X/Y/Z-axis respectively.
 - ♦ Velocity: Along the speed vector of the particle.
 - ♦ Lock: Locks the align axis, keeps this orientation, the billboard aligns only along one axis to itâ← s target
- Tilt: Rotation angle of the billboards planes. A tilt of 1 rotates by 180 degrees (turns the billboard upside down).
- Rand: Random variation of tilt.

The animation of the UV textures is a bit tricky. The UV texture is split into rows and columns (N times N). The texture should be square. You have to use UV Split in the UV channel and fill in the name of the UV layer. This generated UV coordinates for this layer.

- UV Split: The amount of rows/columns in the texture to be used.
- Animate: Dropdown menu, indicating how the split UVs could be animated (changing from particle to particle with time):
 - ♦ None: No animation occurs on the particle itself, the billboard uses one section of the texture in itâ← s lifetime.
 - ♦ Time: The sections of the texture are gone through sequentially in particlesâ← lifetimes.
 - ♦ Angle: Change the section based on the angle of rotation around the Align to axis, if View is used the change is based on the amount of tilt.
- Offset: Specifies how to choose the first part (of all the parts in the nÃ n grid in the texture defined by the UV Split number) for all particles.
 - ♦ None: All particles start from the first part.
 - ♦ Linear: First particle will start from the first part and the last particle will start from the last part, the particles in between will get a part assigned linearly from the first to the last part.
 - ♦ Random: Give a random starting part for every particle.
- OffsetX: Offset the billboard horizontally in relation to the particle center, this does not move the texture.
- OffsetY: Offset the billboard vertically in relation to the particle center.
- OB: The target object that the billboards are facing. By default, the active camera is used.
- UV Channel: Billboards are just square polygons. To texture them in different ways we have to have a

way to set what textures we want for the billboards and how we want them to be mapped to the squares. These can then be set in the texture mapping buttons to set wanted textures for different coordinates. You may use three different UV layers and get three different sets of UV coordinates, which can then be applied to different (or the same) textures.

- Normal: Coordinates are the same for every billboard, and just place the image straight on the square.
- Time-Index (X-Y): Coordinates actually define single points in the texture plane with the x-axis as time and y-axis as the particle index. For example using a horizontal blend texture mapped to color from white to black will give us particles that start off as white and gradually change to black during their lifetime. On the other hand a vertical blend texture mapped to color from white to black will make the first particle to be white and the last particle to be black with the particles in between a shade of gray.
- Split: Coordinates are a single part of the UV Split grid, which is a nÃ n grid over the whole texture. What the part is used for each particle and at what time is determined by the Offset and Animate controls. These can be used to make each billboard unique or to use an â← animatedâ← texture for them by having each frame of the animation in a grid in a big image.

- UV: Set the name of the *UV layer* to use with billboards (you can use a different one for each UV Channel). By default, it is the active UV layer (check the Mesh panel in the Editing context, F9).

The main topic here will be the settings in the Extras panel in the Particle sub-context, but we will also speak about using textures, vertex groups or Ipo curves.

- Particles may interact with other objects (e.g. collide) or be influenced by force fields. You can confine this interaction to a group of objects. Also you can configure the behavior of particles after collisions. Force fields and collision objects have to share a common layer with the particle system.
- You can use Vertex Groups to control a multitude of particle attributes, especially the density of the particles. Zero density means zero particles.
- Size and mass of particles govern their physical behavior. E.g. if you have an object attached to the particle you normally donâ← t want to collide the center of the object, but itâ← s border.
- Particle attributes can be animated, especially material animations are important for a good visualization. An animation of 100 frames in the Ipo window is done in the lifetime of each particle. This is the default setting in Blender and is called *relative time*. You may change this assignment with the Time settings.



Time is difficult:

You will save yourself a lot of headache if you read the section about time carefully. Many problems with the particle system are related to the difficult time concept.

- Many of the particle attributes may not only be controlled with a vertex group, but also by a texture, e.g. the emission time, life time or the density.

Time

A particle has a certain lifetime. It is emitted, lives and dies. If you animate certain particle attributes they do not change in the absolute time of your animation, but for each particle individually during itâ← s lifetime. We have seen examples for this in various tutorials.

Sadly this is not always the case. I will list the cases which are animated that way, and also some of the important cases where absolute time is used.

Relative time is used by:

- The particle Ipos with the exception of E_Freq.
- The material attributes of the particle material.
- Object Ipos (only the **object** Ipos!) for object visualization of particles. You can use object Ipos for individual animations.

Absolute time is used by:

- The E_Freq particle Ipo channel.
- Material Ipos for an object (so you canâ← t change the color of an object in relative time).
- Pose Ipos for armatures (this is really a pity).

The Extras Panel

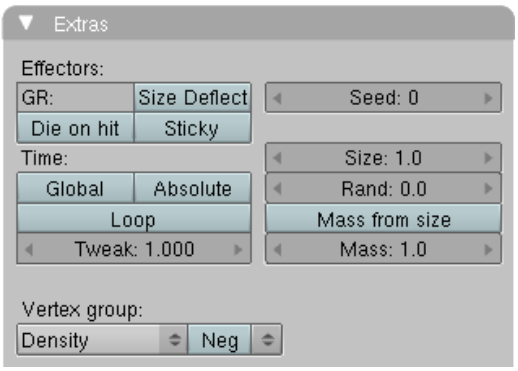


Image 1: The Extras panel in the Particle sub-context.

Effectors

- GR: Limit effectors (force fields and collision objects) to objects in a group.
- Size Deflect: Use particleâ← s size when it deflects from a surface, instead of the center of the particle.
- Die on hit: Particles die when they hit a surface, i.e. stop their movement and disappear. This holds true also for attached objects. If you want to render a particle (or object) after itâ← s death, you have to activate the Died button in the Visualisation panel.
- Sticky: Particles stick to the object they hit, i.e. if the object moves, it carries the sticky particles.

Time

These settings govern the mapping of Ipo curves to particle time. Preset is *relative object time*, 100 frames of Ipo animation are mapped to the lifetime of the particles. The lifetime of the particles is relative to the object time, which is by default the same as the global time. But you can change the object time, e.g. with a Time Ipo.

- Global: Particles are calculated in global time, independent from the object time (but still relative to the particle).
- Absolute: All Ipos that effect particles are calculated in absolute time, and are no longer mapped to the particle time (but may depend from the object time).
- Loop: Particles are reborn once they die.
- Tweak: A multiplier to the physics calculation timestep length, with 1.0 one frame corresponds to 1/25 seconds (for a render done at 25 fps), 2.0 doubles the speed, etc.

Vertex group

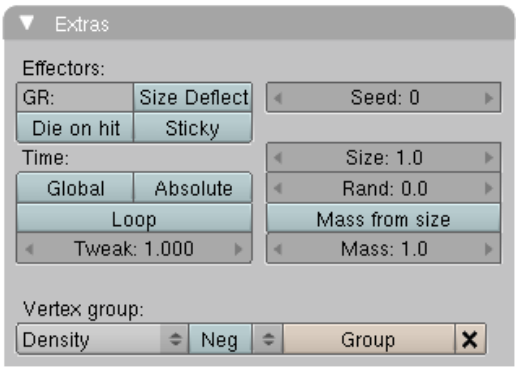


Image 2: Controlling the emission of particles with a vertex group.
With vertex groups you can control e.g. the emission of particles, the Weight of the vertices controls the emission time respectively the strength of the influence. It may be necessary to change the particle system type (Hairâ Emitterâ Hair) to propagate changes in vertex groups.

- Attribute: The attribute of the particles that is effected by the vertex group.
 - ♦ Density: The density of the particles. The number of particles is constant, more particles are emitted from a smaller area.
 - ♦ Velocity: A factor for the speed of the particles.
 - ♦ Length: The path length of Children particles.
 - ♦ Size: Size of the particles.
 - ♦ TanVel: A factor for the Tan velocity during emission.
 - ♦ TanRot: Changes the rotation of the Tan vector.
 - ♦ Effector: Hair particle systems can react directly to force fields, you can limit this influence with a vertex group.
 - ♦ The other parameters are equivalent to the Children parameter settings.
- Neg: Negates the effect of the vertex group.
- Vertex group: Select the vertex group that is used.

Other

- Seed: An offset in the random table that is used for random things in the particles. All random functions in Blender are pseudo random, the randomness is calculated with an unambiguous function. So if you let particles emit with Random, all systems are exactly synchronized and would occupy the same space if their emitter is at same place. With the Seed value you can initialize the random function with different values.

- Size: The size of the particles. Size may seem a bit strange setting since particles are nearly by definition point-like, but the size may be thought of as the particlesâ← area of influence (reactor targets etc.) and the size is also used to scale different visualizations like objects, groups and billboards. Size does not affect the Halo rendering of particles.
- Rand: Gives a random variation to the size.
- Mass from size: Multiply particle mass with its size, so you can create particles with varying mass.
- Mass: Mass is used in physical calculations, a resistance against accelerationâ←¦

If you use a Hair system, you have two additional settings in the Extras panel. Hair can react directly to force fields, even without using softbodies.

- Stiff: The stiffness of Hair systems when a force field is applied.
- Children: Applies force fields to children, and not to the parent particles.

Textures controlling particle attributes

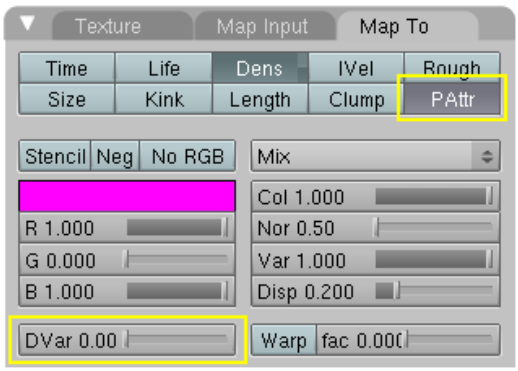


Image 3: Controlling the emission of particles with a texture.
Instead of using vertex groups you can control many of the particles parameters with a texture. The button PAttr (â← particle attributesâ←) only appears in the Map To panel if youâ← ve created a working particle system.

You may use only Orco or UV in the Map Input panel, if you set other input coordinates then Orco will be used silently instead.

The result of the texture is depending on the value of DVar (â← Destination variableâ←). E.g. for density the default value is 1, if you want to adjust this value with a texture you have to set DVar to zero. The notable exception of that rule is the Time parameter, this is directly set with the value of the texture.

- Time: The emission time in relation to the Sta/End values. A texture value of 0 means emission at the beginning, a value of 1 emission at the end.
- Life: Multiplication factor (0 to 1) for the lifetime of the particles. Since 1 is the default value, you can only lower the lifetime with a DVar less than 1.
- Dens: The density of particles. Since particles have a default density of 1, you can only lower the density with a DVar less than 1.
- IVel: Multiplication factor for the initial velocity. This doesnâ← t change the direction, only the value. So if you want to randomize the emission speed but still want the emission in the normal direction, this is what you need.
- Rough: Factor for the roughness of the children.
- Size: Factor for the size of the particles.
- Kink: Factor for the kink frequency of children.

- Length: Factor for the length of the children.
- Clump: Factor for clumping of the children.

Ipo channels

The Ipo type Particles does only appear if the active object carries a particle system.

All Ipo channels with â← E_â←¦â← are working for the emission, the others are working during the lifetime of the particles.

- E_Freq: Emission frequency, the number of particles per frame. The total amount of particles is still set with the parameter Amount, also the Start and End frame is not changed by that. This is the only channel that is independent from the life time of the particles, it is evaluated in absolute time. All other channels are (by default) evaluated in relative time.
- E_Live: Factor for the lifetime of the particles.
- E_Speed: Factor for the emission speed of the particles.
- E_Angular: Factor for the emission rotation of the particles, it needs also Dynamic to be set.
- E_Size: Factor for the emission size.

The following channels change the behavior of the particles during their lifetime.

- Angular: Rotation speed.
- Size: Size.
- Drag/Brown/Damp: According to the parameters in the Physics panel.
- Length: Only for Path visualization (Hair systems or Keyed particles): length of the hair.
- Clump: Only for Path visualization with Children: parameter Clump for the Children.
- GravX/GravY/GravZ: Parameter AccX/AccY/AccZ in the Physics panel.
- KinkAmp/KinkFreq/KinkShape: Only for Path visualization with Children.
- BBTilt: Only for Billboard visualization: tilt of the billboard.
- F-parameters: Regarding forcefields which are produced by particles. The same as for objects.
- f2-parameters: Since particles can produce two force fields, the settings for the second force field.

Controlling particle parameters

An overview about the possibilities to set particle parameters is given in the following table. Setting parameters with a texture means PAttr on in the Map To panel, else the name of the Ipo channel or the Vertex group attribute is stated. A blank field means that you canâ← t influence the parameter by that method.

What can effect what

	Emit Ipo	**Particle Ipo**	**Texture**	**Vertex group**
Emission time	E_Freq		Time	
Emission location			Dens	Density
Lifetime	E_Life		Life	
Starting speed	E_Speed		IVel	Velocity

Angular velocity	E_Angular	Angular		TanVel/TanRot
Size	E_Size	Size	Size	Size
Length		Length	Length	Length
Gravity (Accâ←¦)		GravX/Y/Z		
Brownian force		Brown		
Damping		Damp		
Drag		Drag		
Clump		Clump	Clump	Clump
Kink amplitude		KinkAmp		
Kink frequency		KinkFreq	Kink	Kink
Kink shape		KinkShape		
Rough			Rough	Rough
Billboard tilt		BBTilt		

Children are Hair and Keyed particles assigned subparticles. They make it possible to work primarily with a relatively low amount of Parent particles, for whom the physics are calculated. The children are then aligned to their parents. Without recalculating the physics the number and visualization of the children can be changed.

- Children can be emitted from particles or from faces (with some different options). Emission from Faces has some advantages, especially the distribution is more even on each face (which makes it better suitable for fur and the like). However, children from particles follow their parents better, e.g. if you have a softbody animation and donâ← t want the hair to penetrate the emitting mesh. But see also our manual page about Hair.
- If you turn on children the parents are no longer rendered (which makes sense because the shape of the children may be quite different from that of their parents). If you want to see the parents additionally turn on the Parents button in the Visualization panel.
- Children carry the same material as their parents and are colored according to the exact place from where they are emitted (so all children may have different color or other attributes).

The possible options depend from the type of particle system, and if you work with *Children from faces* or *Children from particles*. We donâ← t show every possible combination, only the settings for a Hair particle system.

Particles

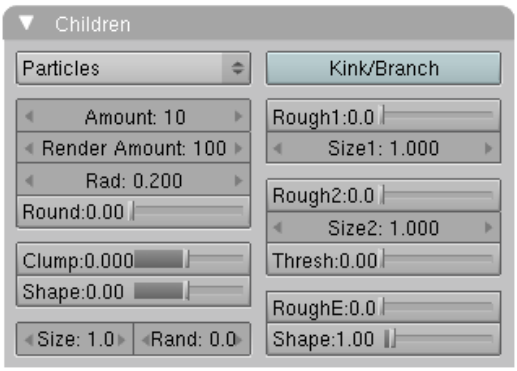


Image 1: The Children panel for *Children from Particles*.

- Amount: The number of children in the 3D window.
- Render Amount: The number of children to be rendered (up to 10.000).
- Rad: The radius in which the children are distributed around their parents. This is 3D, so children may be emitted higher or lower than their parents.
- Round: The roundness of the children around their parents. Either in a sphere (1.0) or in-plane (0.0).
- Clump: Clumping. The children may meet at their tip (1.0) or start together at their root (-1.0).
- Shape: Form of Clumping. Either inverse parabolic (0.99) or exponentially (-0.99).

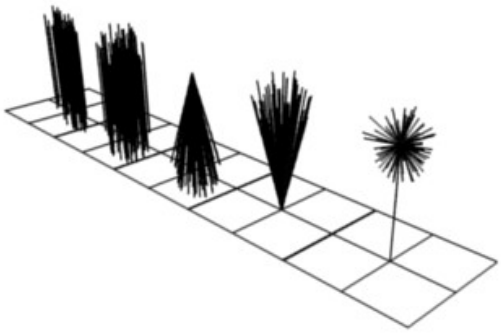


Image 2: From left to right: Round: 0.0 / Round: 1.0 / Clump: 1.0 / Clump: -1.0 / Shape: -0.99.

- Size/Rand: Only for Emitter and Reactor systems. A multiplier for child particle size in relation to the size of their parents, and a random variation.
- Rough1/Size1: It is based on children location so it varies the paths in a similar way when the children are near.
- Rough2/Size2/Thresh: It is based on a random vector so itâ← s not the same for nearby children. The threshold can be specified to apply this to only a part of children. This is useful for creating a few stray children that wonâ← t do what others do.
- RoughE/Shape: â← Rough Endâ← randomizes path ends (a bit like random negative clumping). Shape may be varied from <1 (parabolic) to 10.0 (hyperbolic).

Kink/Branch

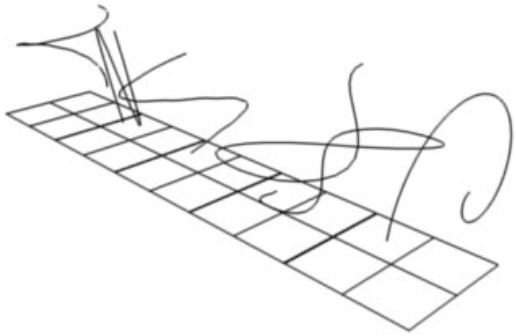


Image 3: Child particles with Kink. From left to right: Curl / Radial / Wave / Braid / Roll.
With Kink you can rotate the children around the parent, Branch branches the hair (dependent on the number of children). These settings donâ← t work alternatively with roughness, but additionally (it seems there would have been another panel necessary for all the options).

- See above picture (*Image 3*) for the different types of Kink.
- X/Y/Z: Axis for the offset of the children.
- Freq: The frequency of the offset (1/total length). The higher the frequency the more rotations are done.
- Shape: Where the rotation starts (offset of rotation).
- Amplitude: The amplitude of the offset.

For branching the parameter Threshold is especially important. 0.0 branches at the very tip of the particles (i.e. no branching at all), 1.0 branches at the root of the particles.

- Animated: Branches in every frame of an animation in a new, random way.
- Symmetric: Start- and end-points are the same.

Faces

Children may also be emitted between the *Parent* particles on the faces of a mesh. They interpolate between adjacent parents.

- This is especially useful for fur, because you can achieve an even distribution.
- Some of the children can become virtual parents, which are influencing other particles nearby.
- *Children from faces* can have *Child simplification*.

Many options are the same as for *Children from particles*, but you donâ← t have roundness and branching. There are two additional settings:

- Spread: Spread children from the faces.
- Use Seams: Use seams to determine parents. This makes it easier to create sharp partings.

Here we will discuss one of the most amazing features of the new particle system: the Particle Mode. In this mode you can edit the keypoints (=controlpoints) of Editable Hair particle systems. The number of keypoints is the number of segments plus one. Additionally you can add Hair particles.

Since working in particle mode is pretty easy and very similar to working with vertices in the 3D window, we will show how to set a particle system up and then give a reference of the various functions.

How to get the Particle Mode up and running

- Create a Hair particle system.
- Give it an initial velocity in Normal direction.
- Click on Set Editable.
- Select Particle Mode in the Mode dropdown-box in the header of the 3D window.

To see what you will be doing:

- Select Point select mode (in the header of the 3D window).
- Turn on the Particle Edit Properties (*PEP*) panel with N.

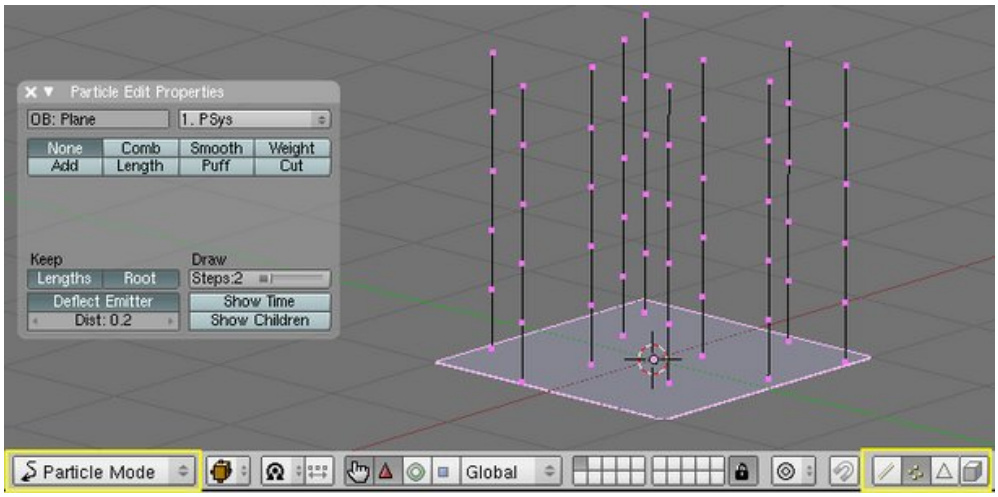


Image 1: Setting up particle mode.
Now your system should look similar to (*Image 1*), probably with more particles. Go ahead and edit the keypoints.

Particle Mode Options

Selecting keypoints

- Single: RMB .
- All: A.
- Linked: Move the mouse over a keypoint and press L.
- Border select: B.
- First/last: W â Select First/Select Last.

You may also use the Select Menu.

Moving keypoints or particles

- To move selected keypoints press G, or use one of the various other methods to grab vertices.
- To move a particle including itâ← s root you have to turn off Keep Root in the Particle Edit Properties (*PEP*) panel.
- You can do many of the things like with vertices, including scaling, rotating and removing (complete particles or single keys).
- You may not duplicate or extrude keys or particles, but you can subdivide particles which adds new keypoints (W â Subdivide/2 NumPad).
- Alternatively you can rekey a particle (W â Rekey/1 NumPad) and choose the number of keys.

How smooth the hair follows the keypoints depends on the number of Draw Steps, which you can set either in the *PEP* or in the Visualisation panel.

Mirroring particles

- If you want to create an X-Axis symmetrical haircut you have to do following steps:
 - ♦ Select all particles with A.
 - ♦ Mirror the particles with CtrlM, or use the Particle â Mirror menu.
 - ♦ Turn on X-Axis Mirror Editing in the Particle menu.

It may happen that after mirroring two particles occupy nearly the same place. Since this would be a waste of memory and rendertime, you can Remove doubles either from the Specials (W) or the Particle menu.

Hiding/Unhiding

- Hiding and unhiding of particles works similar as with vertices in the 3D window. Select one or more keypoints of the particle you want to hide and press H. The particle in fact doesnâ← t vanish, only the keypoints.
- Hidden particles (i.e. particles whose keypoints are hidden) donâ← t react on the various brushes. But:
- If you use Mirror Editing even particles with hidden keypoints may be moved, if their mirrored counterpart is moved.

Select Modes

- Path: No keypoints are visible, you can select/deselect only all particles.
- Point: You see all of the keypoints.
- Tip: You can see and edit (including the brushes) only the tip of the particles, i.e. the last keypoint.

The Particle Edit Properties Panel

The Particle Edit Properties Panel (short *PEP* panel) bears various options to facilitate the working with particle hair.

- With the button row you can select the type of â← Combâ← utility you want to use:
 - ♦ None: No special tool, just edit the keypoints as â← normalâ← vertices.
 - ♦ Comb: Moves the keypoints (similar to â← proportional editingâ←).
 - ♦ Smooth: Parallels visually adjacent segments.

- ♦ Weight: This is especially useful for softbody animations, because the weight defines the softbody Goal. A keypoint with a weight of 1 wonâ ← t move at all, a keypoint with a weight of 0 subjects fully to softbody animation. This value is scaled by the GMin-GMax range of softbody goals.
- ♦ Add: Adds new particles.
- ♦ Length: Scales the segments, so it makes the hair longer or shorter.
- ♦ Puff: Rotates the hair around itâ ← s first keypoint (root). So it makes the hair stand up (Add) or lay down (Sub).
- ♦ Cut: Scales the segments until the last keypoint reaches the brush.

- Keep:
 - ♦ Length: Keep the length of the segments between the keypoints when combing or smoothing the hair. This is done by moving all the other keypoints.
 - ♦ Root: Keep first key unmodified, so you canâ ← t transplant hair.
 - ♦ Deflect Emitter/Dist: Donâ ← t move keypoints through the emitting mesh. Dist is the distance to keep from the Emitter.
- Draw:
 - ♦ Steps: Drawing steps (same as in Visualisation panel).
 - ♦ Show Time: Shows the frame in which the position is reached (quite useless now but very promising for the future).
 - ♦ Show Children: Draws the children of the particles too. This allows to fine tune the particles and see their effects on the result, but it may slow extremely down if you have many children.

If you use an older version of Blender: documentation for version 2.45 and earlier.



Image 1: Particle hair with Softbody animation (blend file).
Animation of *Image 1* (info)
This page does not contain any â ← newâ ← information, it tries to compile methods to create particle hair as realistic as possible. I invite any experienced user to discuss the contents of this page on the discussion page of this article, you may well have other experiences.

There are always several ways to achieve a certain result. Some road marks up front:

- The easiest thing to create is fur. If you have 2GB RAM you may render up to 2.000.000 particles quite fast. If you need more particles, which may happen with several large characters in the foreground, you need more memory or have to split the scene.
- If you want to animate the hair with Soft Body physic, take care of the amount of particles (some 1.000). Else the soft body calculations may take really long time. Collision is facilitated if you use an

invisible, simplified collision mesh instead of the original mesh (no ears, no eye socket).

Collision is still not working perfectly:

Your best way to go is to use it as little as possible. Fake it, weight paint it, post produce it - I have not managed yet to create 100% perfect collision, though you can come pretty close.

- You have to use weight painting if you want to animate with soft body physics.
- Children can be emitted from Faces or from Particles. Children from faces have some advantages:
 - They are distributed more even on the mesh, even if you have a relatively small amount of parents.
 - They always start from the mesh.
 - You can Use Seams to part the particles.
 - You have a *Level of Detail* function, i.e. the objects will be rendered simpler when the distance to the camera increases. So you can use more particles.
- Children from particles do follow their parents better though.

If you have not yet read the fur tutorial, it would be a good idea to do it now and come back after that.

Particle system



Image 2a: A first, simple particle system.

- Weâ← re going to start with an UVSphere as â← headâ← , activate AutoSmooth and then SetSmooth.
- Add a Subsurf modifier.
- Add a new particle system. 1000 particles are enough here, so we can work fast in the 3D window. For a real hair cut I would use more particles (5.000 upwards), but to animate them with soft body physics is quite demanding for the CPU.
 - Type Hair.
 - Emit from Faces.
 - Activate Random and Even, so that the hair distributes evenly on the object. Without Even small faces would emit as many particles as large faces.
- Set the parameter Normal in the Physics panel to 0.5.

The hair has now a certain length, the emitter is invisible in the render (*Image 2a*).

To let the particles be emitted only from a certain part of the mesh there are three possibilities:

1. Use a vertex group.
2. Use a texture.
3. Use another mesh.

Iâ← ve used a vertex group.

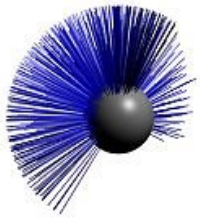


Image 2b: The emission was controlled with a vertex group. The emitter uses another material than the particles.

- Create a vertex group in Edit mode. The simplest way to do that would be to select the regarding vertices and press CtrlG â Add selected to new Group. Rename that group to â← `Hair`â← .
- Use this group in the Extras panel in the section Vertex Group. Attribute: Density. If the hair is still emitted from the complete object change the particle system type to Emitter and back to Hair, now it should work.
- To render the emitter also, activate Emitter in the Visualization panel.
- Use material no.2, we donâ← t have any materials till now, but weâ← re going to create some soon.
- Since were going to work with many particles, we use the button Strand Render in the Visualization panel. This activates the keypoint strand render. See the manual page regarding Strands for pros and cons of the different strand shaders.
- Create two new materials for the object in the Editing context, the Link and Materials panel, and assign them both. The material that shall be used for the emitter object should be assigned at last.
- It is really a good habit to name the materials properly. Iâ← ve named the first material â← `Skin`â← , the second material â← `Hair`â← . The â← `Hair`â← gets a simple, blue material (*Image 2b*).

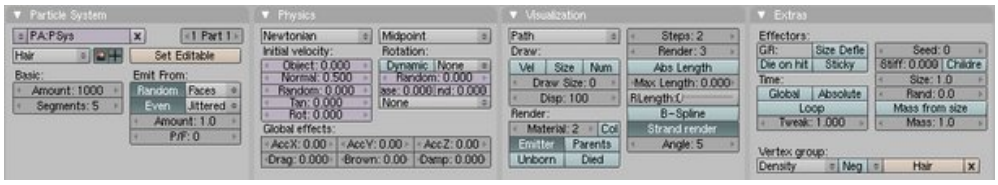


Image 2c: Settings for the particle system till now.

Add volume and dress

To increase the amount of hairs, we add *Children*.



Image 3a: 10 children and a bit dressed.

- Activate Children from: Faces in the Children panel. Set the Render Amount to 10, so that you can do

relatively fast test renders.

Children from: Faces donâ← t follow their parents too well. So probably some of the children will penetrate the head (e.g. at the ears). This problem canâ← t be avoided automatically. If you have errand hairs visible in a still image the fasted way may be to get rid of them with post processing. Children from: Particles do follow their parents better, especially if you use many parents.

- Click on Set Editable in the Particle System panel to dress the hair.
- Change to Particle mode with the dropdown list in the window header of the 3D window.

Now by default only the parents are visible, because we can edit only them.

- N activates the Particle Edit Properties panel.
- You see the control points of the hair in Point Select mode (directly next to Limit Selection to Visible in the window header, see Particle Mode).

Now dress the hair to your liking. This may take itâ← s time, Iâ← ve found it took quite a few attempts to get any useful hairstyle (*Image 3a*). If you have selected single control points you edit only the selection. You can hide controlpoints with H.

If you want to animate the hairs as softbodies, you have to adjust the weight of the controlpoints. Points with a weight of 0 (in black) follow the softbody physics completely, points with a weight of 1 (white) donâ← t follow the softbody animation. Hairs directly adjacent to the head should get a weight of 1, so that it is not necessary to calculate the collision. I would painstakingly difficult to use softbody animation for hairstrands behind the ear, this will never work correctly.

Use a small weight for the tips of long hair and the sections that shall move freely.

Insert a basic lighting (also a Spot lamp with shadow buffer), to judge the flow of the hair and the material. If the hair is too angled (like in *Image 3a*), increase the amount of Rendersteps (Visualization panel) carefully (only one at a time). Many rendersteps need much RAM. Each segment is subdivided into smaller rendersegments.

Material

To let the hair look like hair, you have to change the hair color and itâ← s thickness. At first we will set the hair color with the RGB color of the material. You can also create color variations with a texture. The basis color of the particle is the color of itâ← s emission point. Often we use an additional texture along the strand to fade the hair towards its tip. That makes the hair more fluffy looking and softer. At last you need the right amount of children. Number of particles and thickness of hair create the effect youâ← re looking for. If you have 1.000 parent particles you need around 50 to 100 children. Thin hair (blond) need more, thick hair (black) need less particles.

The color of the hair is - especially with thin hair - governed by the specular highlights, so it may be wise to set the Spec color also.

Black hair



Image 4a: 140.000 particles, keypoint strands.
Black hair is quite simple, so we will begin with that. Simply set the RGB color to black.

- Strand-Shader: The most important settings are in the Materials tab, hidden behind the button Strands in the Links and Pipeline panel. Hair stays nearly as thick at the end as at the beginning. Start: 0.86, End: 0.854, Shape: -0.686.
- Textures: We need at least one texture, to fade the strand towards the tip. Map Input: Strand, Map To: Alpha, DVar: 0. With a custom blend texture we can fade the strand at discretion.

To make the hair really transparent, you have to activate ZTransp.

Blond hair

- Strand settings:
 - ♦ This time we use Blender Units. Set Start to 0.003 and End to 0.002 (if you need thinner hair you have to enlarge the object). Set Minimum to 2. This fades the hair, without using an additional texture.
 - ♦ Leave Shape on 0, or set it to something like -0.5. The hair will then become relatively early (at its root) thinner, but the hair is already so thin, that it may be to much.
- Material settings:
 - ♦ The hair color is relatively dark (0.37, 0.30, 0.22). The hair will become blond because the specular highlights are bright.
- Textures: With a texture you change the color of the specular highlights. Map To panel settings:
 - ♦ ◊ Csp.
 ◊ Color (0, 0, 0) (black).
 ◊ Use a relatively high resolution, random texture, for example a Marble texture with a high Turbulence value.

The material I have used for (*Image 1*) is by courtesy of Len. Itâ← s the same that he has used for his Saskia.

Animation

For animating hair particles with the Soft Body simulation see Combination of Soft Bodies with hair particles.

Links

- Blender Material Repository with hair material

If you are using a version of Blender up to 2.46, click here to access the old documentation.



Image 1a: A softbody cloth uncovering a text. Animation â← Blend file
A Soft Body is a physical simulation of what would happen to an object if it were to exist â← in the real worldâ← . This is done by applying forces to the vertices or controlpoints of the object. There are exterior forces like gravity or forcefields and interior forces that hold the vertices together. This way you can simulate the shapes that an object would take on in reality if it had volume, was filled with something, and was acted on by real forces.

Soft Bodies can interact with other objects (*Collision*). They can interact with themselves (*Self Collision*).

The result of the Soft Body simulation can be converted to a static object. You can also *bake edit* the simulation, i.e. edit intermediate results and run the simulation from there.

Typical scenarios for using Soft Bodies



Image 1b: A wind cone. The cone is a Soft Body, as the suspension. Animation â← Blend file
Soft Bodies are well suited for:

- Elastic objects with or without collision.
- Flags, fabric reacting to forces.
- Certain modeling tasks, like a cushion or a table cloth over an object.
- Blender has another simulation system for clothing (see Clothes). But you can sometimes use Soft Bodies for certain parts of clothing, like wide sleeves.
- Hair (as long as you minimize collision).
- Animation of swinging ropes, chains and the like.

The following videos may give you some more ideas: [1], [2]

Creating a Soft Body

Mode: Object Mode

Panel: Object context â Physics sub-context â Soft Body, and Editing context â Modifiers

Hotkey: F7 to get to Object context; repeat to change sub-context.

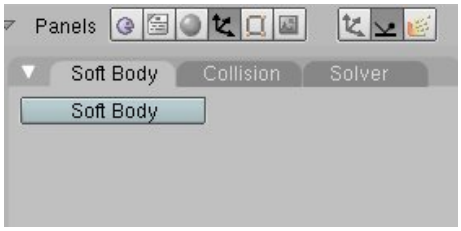


Image 2a: Create a Soft Body.



Image 2b: Create a Soft Body from a Hair particle system.

Soft Body simulation works for all objects that have vertices or control points:

- Meshes.
- Curves.
- Surfaces.
- Lattices.
- Hair particles.

To activate the Soft Body simulation for an object:

- Change to the Object context (F7), and select Physics sub-context.
- Activate the Soft Body button in the Soft Body panel (*Image 2a*).
- For a hair particle system first select the particle system, than activate Soft Body (*Image 2b*).

A lot of options appear. For a reference of all the settings see this page.

- You start a Soft Body simulation with AltA.
- You pause the simulation with Space, continue with AltA.
- You stop the simulation with Esc.

Cache and Bake

Soft Bodies use a unified system for caching and baking (together with cloth and particles). The results of the simulation are automatically cached to disk when the animation is played, so that the next time it runs, it can play again quickly by reading in the results from the disk. If you Bake the simulation the cache is protected and you will be asked when youâ← re trying to change a setting that will make a recalculating necessary.



The simulation is only calculated for the frames in-between the Start and End frames (Bake panel), even if you donâ← t actually bake the simulation! So if you want a simulation longer than the default setting of 250 frames you have the change the End frame.:

{{{2}}}

- Caching:
 - ♦ As animation is played, each physics system writes each frame to disk, between the simulation start and end frames. These files are stored in folders with prefix â← `blendcache`â← , next to the .blend file.
 - ♦ The cache is cleared automatically on changes - but not on all changes, so it may be necessary to free it manually, e.g. if you change a force field. Note that for the cache to fill up, one has to start playback before or on the frame that the simulation starts.
 - ♦ If you are not allowed to write to the required sub-directory caching will not take place.
 - ♦ The cache can be freed per physics system with a button in the panels, or with the CtrlB shortcut key to free it for all selected objects.
 - ♦ You may run into trouble if your .blend file path is very long and your operating system has a limit on the path length that is supported.
- Baking:
 - ♦ The system is protected against changes after baking.
 - ♦ The Bake result is cleared also with CtrlB for all selected objects or click on Free Bake for the current Soft Body system.
 - ♦ If the mesh changes the simulation is not calculated anew.

For renderfarms, it is best to bake all the physics systems, and then copy the blendcache to the renderfarm as well.

Interaction in real time

To work with a Soft Body simulation you will find it handy to use the Timeline window. You can change between frames and the simulation will always be shown in the actual state. The option Continue Physics in the Playback menu of the Timeline window lets you interact in real time with the simulation, e.g. by moving collision objects or shake a Soft Body object. And this is real fun!



Continue Physics does not work while playing the animation with AltA:

Right. This works only if you start the animation with the Play button of the Timeline window.

You can than select the Soft Body object while running the simulation and Apply the modifier in the Modifiers panel of the Editing context. This makes the deformation permanent.

Tips

- Soft Bodies work especially well if the objects have an even vertex distribution. You need enough vertices for good collisions. You change the deformation (the stiffness) if you add more vertices in a certain region (see the animation of *Image 1b*).
- The calculation of collisions may take a long time. If something is not visible, why calculate it?
- To speed up the collision calculation it is often useful to collide with an additional, simpler, invisible, somewhat larger object (see the example to *Image 1a*).
- Use Soft Bodies only where it makes sense. If you try to cover a body mesh with a tight piece of cloth and animate solely with Soft Body, you will have no success. Self collision of Soft Body hair may be activated, but that is a path that you have to wander alone. We will deal with Collisions in detail later.
- Try and use a Lattice or a Curve Guide Soft Body instead of the object itself. This may be magnitudes faster.

Links

- Developer Notes
- Swinging of a chain
- Softbodies for Rigged Characters

Forces are applied to the vertices (and nearly exclusively to the vertices) of Soft Body objects. This is done using Newtons Laws of Physics:

1. If there is no force on a vertex, it stays either unmoved or moves with constant speed in a straight line.
2. The acceleration of a vertex depends on its mass and the force. The heavier the mass of a vertex the slower the acceleration. The larger the force the greater the acceleration.
3. For every action there is an equal and opposite reaction.

Well, this is done only in the range of computing accurateness, there is always a little damping to avoid overshoot of the calculation.

Exterior Forces

We will begin with a very simple example - the default cube.

- To judge the effect of the external forces you should at first turn off the Goal, so that the vertices are not retracted to their original position.
- Press AltA.

What happens? The cube moves in negative Z-direction. Each of itâ← s eight vertices is affected by a global, constant force - the gravitation. Gravitation without friction is independent from the weight of an object, so each object you would use as a Soft Body here would fall with the same acceleration. The object does not deform, because every vertex moves with the same speed in the same direction.

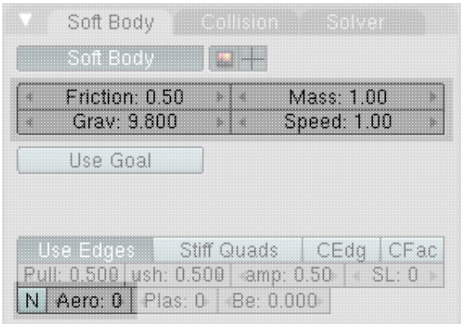


Image 1: Settings regarding external forces for Soft Bodies.
We will discuss the settings here in detail:

Friction
: The friction of the surrounding medium. The larger the friction, the more viscous is the medium. Friction always appears when a vertex moves relative to itâ← s surround medium. If you Use Edges, also Aero can be used for friction.

Grav
: Gravity, the amount of force in Z-axis direction. Earth gravity is a value close to 9.8 $m.s^{-2}$. Positive values make Soft Bodies drop; negative values make them rise.

Mass
: Mass value for vertices. Larger mass slows down acceleration, except for gravity where the motion is constant regardless of mass. Larger mass means larger inertia, so also braking a Soft Body is more difficult.

Speed
: You can control the internal timing of the Softbody system with this value. It sets the correlation between frame rate and tempo of the simulation. A free falling body should cover a distance of about five meters after one second. You can adjust the scale of your scene and your simulation with this correlation. If you render with 25 frames per second and 1 meter shall be 1 BU, you have to set Speed to 1.3.

Aero (Aerodynamik)
: This is special as here (and only here) an exterior force is not applied to the vertices but to the connecting edges. Technically, a force perpendicular to the edge is applied. The force scales with the projection of the relative speed on the edge (dot product). Note that the force is the same if wind is blowing or if you drag the edge through the air with the same speed. That means that an edge moving in its own direction feels no force, and an edge moving perpendicular to its own direction feels maximum force. Try a value of 30 at first.
N uses another calculation model for the aerodynamic force, which creates turbulence. So to get a nice movement of e.g. a flag in the wind, you no longer have to animate the wind force.

To create other forces you have to use another object, often Empty objects are used for that. You can use some of the forces on Soft Body vertices as on Particles. Soft Bodies react only to:

- Spherical
- Wind
- Vortex

Soft bodies do react on Harmonic fields, but not in a useful way. So if you use a Harmonic field for particles move the Soft body to another layer.

See the section Force Fields for details. The force fields are quite strong, a Spherical field with a strength of -1.0 has the same effect that gravity has - approximately - a force of 10 Newton.

Keeping the vertices at their place (â← Goalâ←)

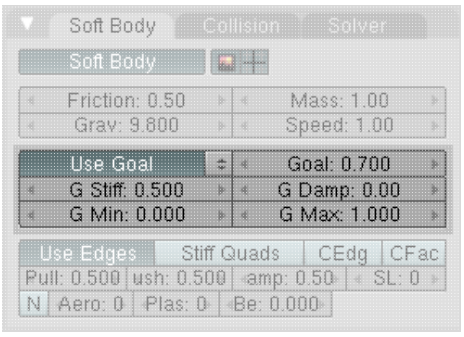


Image 2a: Settings for the Soft Body Goal.

What is a goal? A goal is something that you are trying to reach - that is at least the meaning here for the Soft Body vertices.

You have to confine the movement of vertices in certain parts of the mesh, e.g. to attach a Soft Body object at other objects. This is done with the Goal (target). The target position is the original position of the vertex, like it would result from the â← normalâ← animation of an object including Shape Keys, Hooks and Armatures. The vertex tries to reach itâ← s target position with a certain, adjustable intensity.

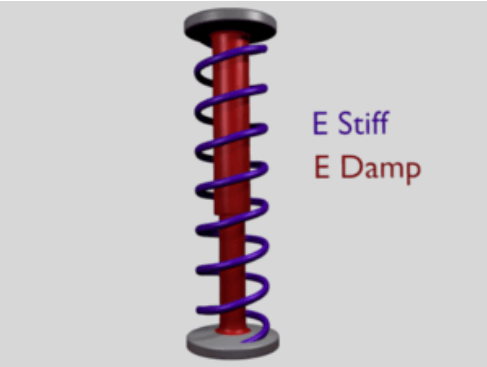


Image 2b: Shock absorber description.
Imagine the vertex is connected with itâ← s target through a spring (*Image 2b*).

Goal

This parameter defines how strong the influence of this spring is. A strength of 1 means, that the vertex will not move as Soft Body at all, instead keep its original position. 0 Goal (or no Goal) means, that the vertex moves only according to Soft Body simulation. If no vertex group is used/assigned, this numeric field is the default goal weight for all vertices. If a vertex group is present and assigned, this button instead shows an popup selector button that allows you to choose the name of the goal vertex group. If you use a vertex group the weight of a vertex defines its goal.
Often weight painting is used to adjust the weight comfortably. For non-mesh objects the Weight parameter of their vertices/controlpoints is used instead (W in Edit mode, or use the Transform Properties panel). The weight of Hair particles can also be painted in Particle mode.

G Stiff

The spring stiffness for Goal. A low value creates very weak springs (more flexible â← attachmentâ← to the goal), a high value creates a strong spring (a stiffer â← attachmentâ← to the goal).

G Damp

The friction of the spring. With a high value the movement will soon come to an end (little jiggle).

G Min/G Max

When you paint the values in vertex-groups (using WeightPaint mode), you can use the G Min and G Max to fine-tune (clamp) the weight values. The lowest vertex-weight (blue) will become G Min, the highest value (red) becomes G Max (please note that the blue-red color scale may be altered by User Preferences).

For now all is applied to single vertices:

For now we have discussed vertex movement independent of each other, similar to particles. Every object without Goal would collapse completely if a non uniform force is applied. Now we will move to the next step, the forces that keep the structure of the object and make the Soft Body to a real Body.

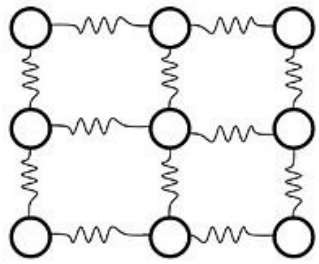


Image 1a: Vertices and forces along their connection edges.
To create a connection between the vertices of a Soft Body object there have to be forces that hold the vertices together. These forces are effective along the edges in a mesh, the connections between the vertices. The forces act like a spring. (*Image 1a*) illustrates how a 3Ã 3 grid of vertices (a mesh plane in Blender) are connected in a Soft Body simulation.

But two vertices could freely rotate if you donâ← t create additional edges between them. Have you ever tried building a storage shelf out of 4 planks alone? Well - donâ← t do it, it will not be stable. The logical method to keep a body from collapsing would be to create additional edges between the vertices. This works pretty well, but would change your mesh topology drastically.

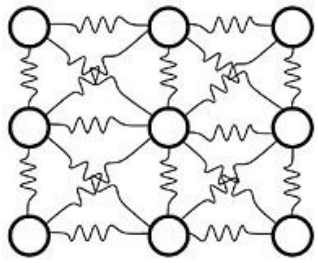


Image 1b: Additional forces with Stiff Quads enabled.
Luckily, Blender allows us to define additional *virtual* connections. On one hand we can define virtual connections between the diagonal edges of a quad face (Stiff Quads, *Image 1b*), on the other hand we can define virtual connections between a vertex and any vertices connected to itâ← s neighbours (Bending Stiffness). In other words, the amount of bend that is allowed between a vertex and any other vertex that is separated by two edge connections.

Edges

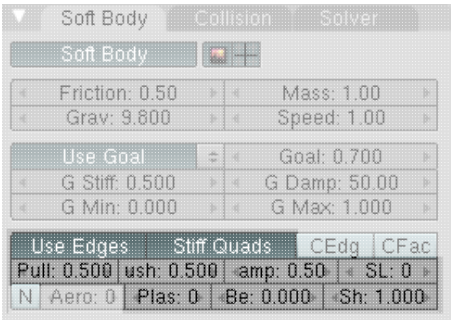


Image 2: Settings regarding the forces between vertices.
The characteristics of a material are set with the Edge properties.

Use Edges
: Allow the edges in a Mesh Object to act like springs.

Pull
: The spring stiffness for edges (how much the edges are allowed to stretch). A low value means very weak springs (a very elastic material), a high value is a strong spring (a stiffer material) that resists being pulled apart. 0.5 is latex, 0.9 is like a sweater, 0.999 is a highly-starched napkin or leather. The Soft Body simulation tends to get unstable if you use a value of 0.999, so you should lower this value a bit if that happens.

Push
: How much the Softbody resist being scrunched together, like a compression spring. Low values for fabric, high values for inflated objects and stiff material.

Damp
: The friction for edge springs. High values (max of 50) dampen the Push/Pull effect and calm down the cloth.

SL (Shrink/Blow)
: The edges can shrink or been blown up. This value is given in percent, 0 disables this function. 100% means no change, the body keeps 100% of his size.

Plas (Plasticity)
: Permanent deformation of the object after a collision. The vertices take a new position without applying the modifier.

Be (Bending Stiffness)
: This option creates virtual connections between a vertex and the vertices connected to itâ← s neighbours. This includes diagonal edges. Damping also applies to these connections.

Stiff Quads
: For quad faces, the diagonal edges are used as springs. This stops quad faces to collapse completely on collisions (what they would do otherwise).

Sh (Shear stiffness)
: Stiffness of the virtual springs created for quad faces.

Stopping an object from collapsing

To show the effect of the different edge settings we will use two cubes (blue: only quads, red: only tris) and let them fall without any goal onto a plane (how to set up collision is shown on the page Collisions).

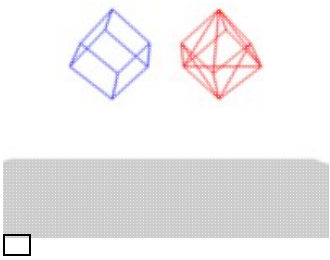


Image 3a: Frame 1 without Stiff Quads.

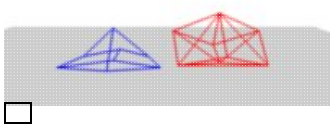


Image 3b: Frame 36.

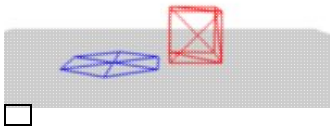


Image 3c: Frame 401.

In (*Image 3*), the default settings are used (without Stiff Quads). The â← quad onlyâ← cube will collapse completely, the cube composed of tris keeps itâ← s shape, though it will deform temporarily because of the forces created during collision.

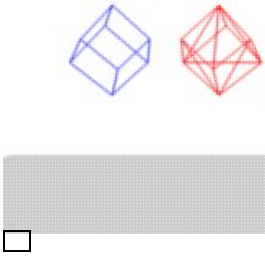


Image 4a: Frame 1 with Stiff Quads.

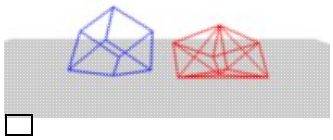


Image 4b: Frame 36.

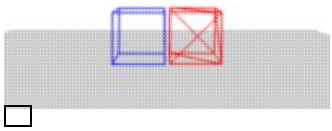


Image 4c: Frame 401.

In (*Image 4*), Stiff Quads is activated (for both cubes). Both cubes keep their shape, there is no difference for the red cube, because it has no quads anyway.

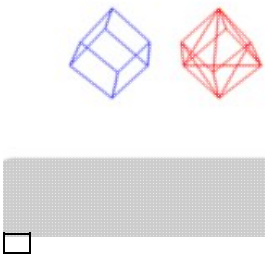


Image 5a: Frame 1 with Bending Stiffness.
Blend file

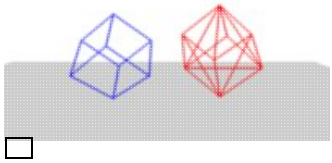


Image 5b: Frame 36.

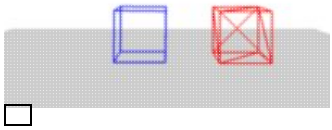


Image 5c: Frame 401.

The second method to stop an object from collapsing is to change itâ ← s Bending Stiffness. This includes the diagonal edges (Damping also applies to these connections).

In (*Image 5*), Be is activated with a strength setting of 1. Now both cubes are more rigid.

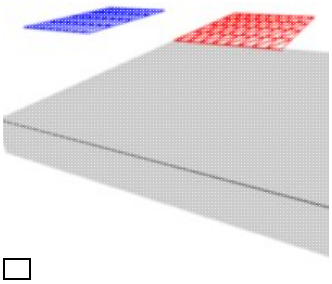


Image 6a: Two planes going to collide.

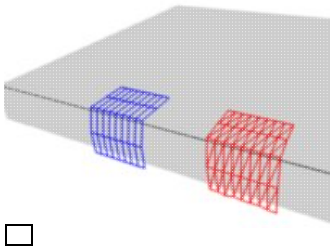


Image 6b: No bending stiffness, Frame 101.

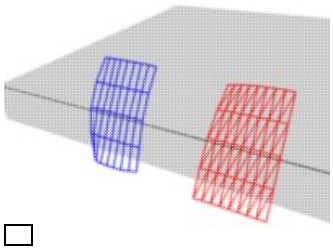


Image 6c: High bending stiffness (10), Frame 101.

Bending stiffness can also be used if you want to make a subdivided plane more plank like. Without Be the faces can freely rotate against each other like hinges (*Image 6b*). There would be no change in the simulation if you activated Stiff Quads, because the faces are not deformed at all in this example.

Bending stiffness on the other hand prevents the plane from being - well - bent.

Attribution

I took some text not only from the original manual page, but also from Basic Physics and Math from Felipeboralli.

--Soylentgreen 07:25, 10 April 2009 (UTC)

There are two different collision types that you may use: collision between different objects and internal collision. We should set one thing straight from the start: the primary targets of the collision calculation are the vertices of a Soft Body. So if you have too few vertices too few collision takes place. Secondarily, you can

use edges and faces to improve the collision calculation.

Collisions with other objects

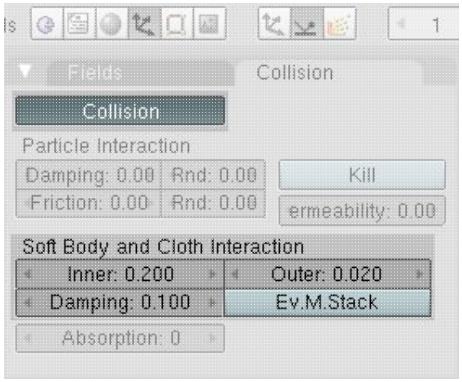


Image 1: Settings for a collision object.
For a *Soft Body* to collide with another object there are a few prerequisites:

1. Both objects have to share a layer, but the layer does not necessarily have to be visible.
2. The collision object has to be a mesh object.
3. You have to activate the option Collision in the Collision panel of the Physics sub-context (*Image 1*) for the collision object. The collision object may also be a Soft Body.
4. If you use modifiers such as Array and Mirror you have to activate EV.M.Stack to ensure that collision calculation is based on the modified object. The sequence of Modifiers is not important.

Motivating examples

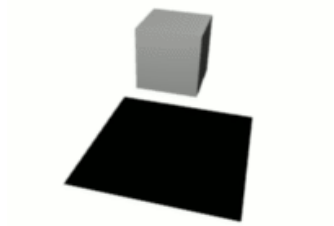


Image 2a: A Soft Body cube colliding with a plane.

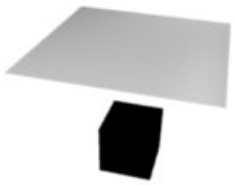


Image 2b: A Soft Body plane colliding with a cube - no interaction at all.

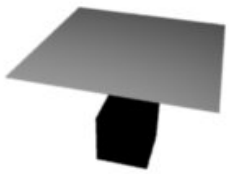


Image 2c: Collision with CFace activated.

A cube colliding with a plane works pretty well (*Image 2a*), but a plane falls right through a cube that it is supposed to collide with (*Image 2b*). Why is that? Because the default method of calculation only checks to see if the four vertices of the plane collides with the cube as the plane is pulled down by gravity. You can activate CFace to enable collision between the face of the plane and the object instead (*Image 2c*), but this type of calculation takes much longer.

Letâ← s have a closer look at the collision calculation, so you can get an idea of how we might optimize it.

Calculating the collisions

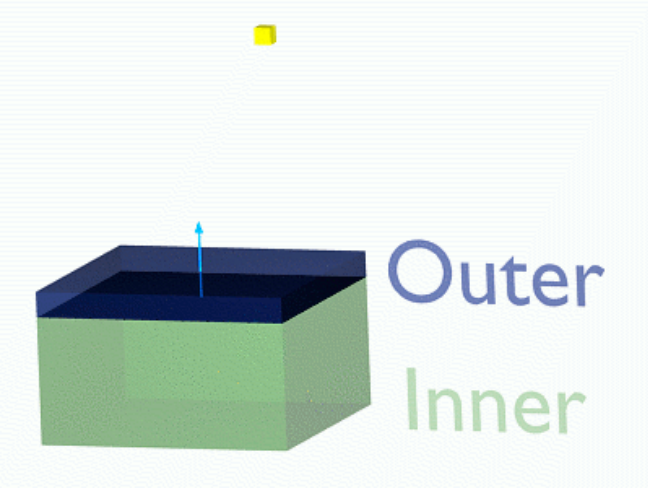


Image 3a: Visualization of the collision of a Soft Body vertex with a plane.

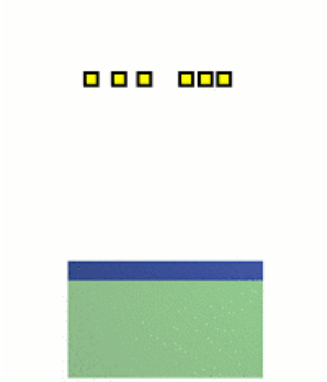


Image 3b: Six Soft Body vertices with different speed.
Blend file

Soft Body simulation is by default done on a per vertex basis. If the vertices of the Soft Body do not collide with the collision object there will be no interaction between the two objects.

In (*Image 3a*), you can see a vertex colliding with a plane. If a vertex penetrates the zone between Outer and Inner, it is repulsed by a force in the direction of the face normal. The position that a vertex finally ends up in is dependent on the forces that act upon it. In the example gravity and the repulsion force of the face balance out. The speed at which the vertex is pulled out of the collision zone is influenced by the Choke parameter (*Image 4*).

Now lets see what happens if we make vertices heavier and let them travel at a faster speed. In (*Image 3b*), you can see vertices traveling at different speeds. The two on the far right (5 and 6) are traveling so fast that they pass right through the collision zone (this is because of the default solver precision - which we can fix later). You will notice that the fourth vertex also travels quite fast and because it is heavier it breaches the inner zone. The first three vertices collide OK.

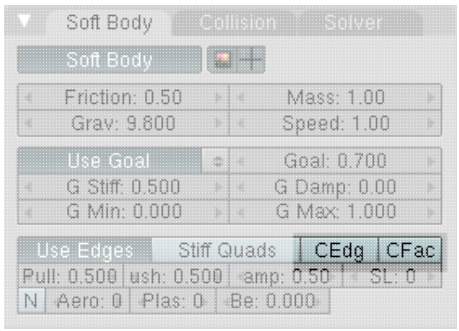


Image 3d: Also *Edges* and *Faces* can be used for the collision calculation.
You can set up your collision so that edges and even faces are included in the collision calculation (*Image 3d*). The collision is then calculated differently. It is checked whether the edge or face intersects with the collision object, the collision zones are not used.

Good collisions

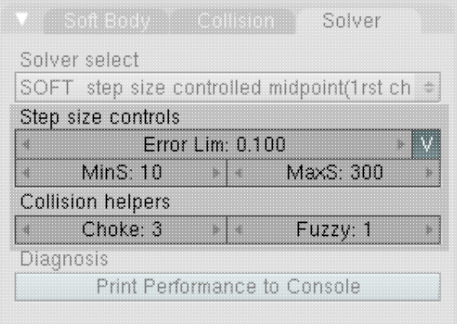


Image 4: Parameters for Soft Body calculation.
If the collision you have set up is not behaving properly, you can try the following:

The best way:

Add Loop Cuts to your Soft Body object in strategic areas that you know are most likely to be involved in a collision.

- The Soft Body object must have more subdivisions than the collision object.
- Check the direction of the face normals.
- If the collision object has sharp spikes they might penetrate the Soft Body.
- The resolution of the solver must match the speed at which Soft Body vertices are traveling. Lower the parameter Error Lim and carefully increase Min S.
- Outer and Inner should be large enough, but zones of opposite faces should not overlap, or you have forces in opposite directions.
- If you use strong forces you should use large zones.
- Set Choke to a high enough value (all the way up if necessary) if you have difficulties with repelled vertices.
- Colliding faces are difficult to control and need long calculation times. Try not to use them.

Often it is better to create a simplified mesh to use as your collision object, however this may be difficult if you are using an animated mesh.

Self Collision

Well, this works pretty well if you simply turn it on. See the Reference page for a detailed explanation of the parameters.

Balls

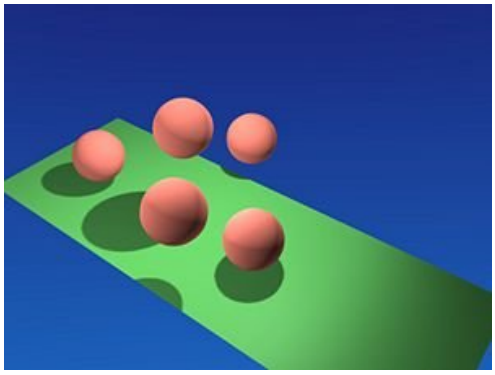


Image 1a: Simple Soft Body balls.
Animation
Blend file
A Soft Body ball is best created from an IcoSphere, because an IcoSphere has evenly spaced vertices and is composed of triangles.

- Insert the standard IcoSphere mesh object.
- Move it a few units in a positive Z direction (upwards).
- Activate Soft Body.
 - The ball is pretty light, every vertex weighs 0.001kg, so the total object weight is 0.042kg.
 - Turn off Goal. We want the ball to be moving freely and reacting only to external forces.
 - Set both Push and Pull to 0.8. Edges tend to be stiffer when they are being pushed together or pulled apart.
 - Set Be(nd) to 0.8. This introduces additional virtual connections between the vertices in the object which restricts the amount of bend (see Interior forces). If you want a harder ball increase Be.
- Activate Collision in the Collision panel. This is not necessary for the ball to collide with other objects, but it is necessary so that other Soft Bodies can collide with the ball.
- Insert a plane and activate Collision on it.
- Ensure that you are at frame 1 of the animation and press AltA.

Now the ball will move in a negative z direction (fall downwards). When it hits the plane it should be repelled.

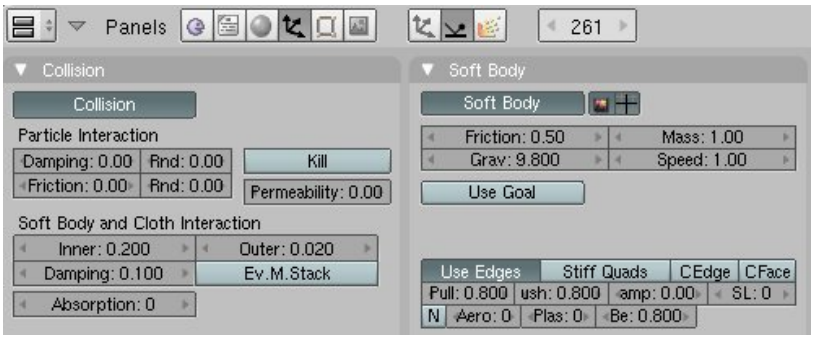


Image 1b: Settings for the Soft Body balls.

Flag in the wind



Image 2a: Soft Body flag in the wind.
Animation

Blend file
This example shows how you can create a simple flag moving in wind (*Image 2a*). The steps are:

1. Create the Flag as a mesh.
2. Assign Weights to a Vertex Group.
3. Designate the flag as a Softbody.
4. Add Wind.

Letâ← s start:

- Remove the cube and save your file.
- Add a plane.
- Rotate it by 90Â° along the X axis (R-X-90).
- Scale it by a factor of 2 in the X-direction.
- Subdivide the plane 4 times. This creates a lot of vertices, but you need that to create nice, smooth movement.
- Add a SubSurf modifier. Unfortunately the simulation will not use the *virtual* vertices that are created by this modifier (at least not in this version of Blender), but it will make the flag simulation look smoother.
- In the Editing context, Link and Materials panel, press Set Smooth.

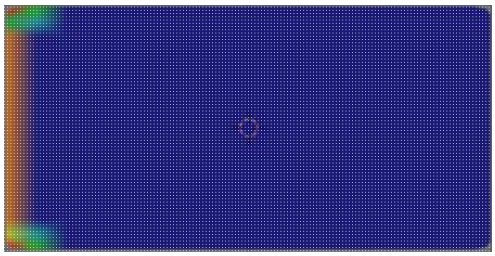


Image 2b: Weight paint for the Soft Body flag.
Now we will use weight painting to pin the leading edge of the flag simulating where the flag would be attached to a flagpole (otherwise it will blow away).

- Change to weight paint mode (see Weight Paint).
- Paint the vertex weights as seen in (*Image 2b*). The corners get a weight of 1. The leading edge of the flag must get a weight close to (but not quite) 1. This is how you go about introducing stiffer zones in a mesh with different weights.

A vertex group called â← `Group`â← is automatically created in this case.

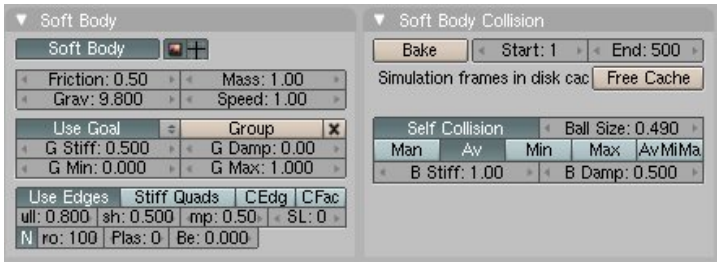


Image 2c: Soft Body settings for the flag.

- Change to Object mode.
- Activate Soft Body.
 - ♦ Select the vertex group â← `Group`â← next to Use Goal. Now vertices with a weight of 1 will not be affected by the simulation, while vertices with a weight of 0 are fully affected by

the simulation.
- ♦ Set Pull to 0.8, so that the flag is not stretched too much.
- ♦ Set Aero to 100, activate N. This will add a bit of random movement to the flag, even without animating the force field.
- ♦ Set the end frame for Bake to 500, because we want the simulation to last longer than 250 frames (even if you donâ ← t bake you still have to adjust the end frame).
- ♦ Activate Self Collision.

- Now, you can press AltA to see the flag react to gravity and air resistance until it comes to rest.

[[Image:Manual-Part-X-SoftBodyExampleWindSetup.png thumb|320px|Image 2d: Example Wind setup.]]

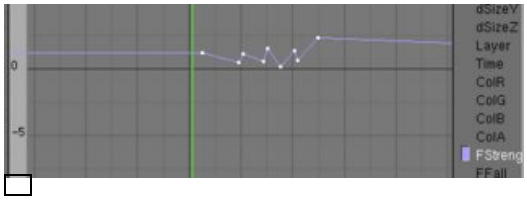


Image 2e: Example Wind Strength IPO.

Now letâ ← s add some wind to the simulation:

- Add an Empty to the scene at the location where we want the wind to blow from. Select Wind from the Fields panel. Set the strength to 0.2.
- Rotate and move the empty so that the Z axis points towards the flag (*Image 2d*). You can temporarily increase the Strength value so that you can more easily see the effect of wind.

Press AltA to see the flag react to wind.

- Add an IPO curve to simulate changing wind strength. This will add more realism to the animation (*Image 2e*).

Interaction

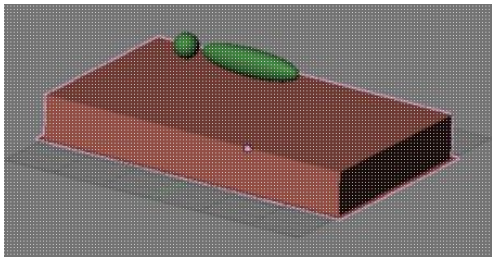


Image 3a: Setup. The green object shall sink into the pink mattress.

Blend file

Weâ ← re going to use the Soft Body simulation as a modeling tool in this example. The goal is to place an object on a mattress and let it sink in a bit.

- Activate collision on the object that will be placed on the mattress.
- Activate Soft Body on the mattress. Make sure the Soft Body has enough subdivisions, to that it will deform properly.

We donâ← t do anything fancy here, so you should stick to the default settings.

- Open a Timeline window.
- In the Playback menu of the Timeline window, select Continue Physics.
- Click on Play Timeline in the window header of the Timeline window.

This is important, donâ← t use AltA to play the simulation here.

- Select the collision object and move it slowly towards the mattress.

The moment the collision object touches the mattress, the mattress will start deforming on collision.

- Move the collision object until youâ← re satisfied with the result.

Donâ← t stop timeline playback!

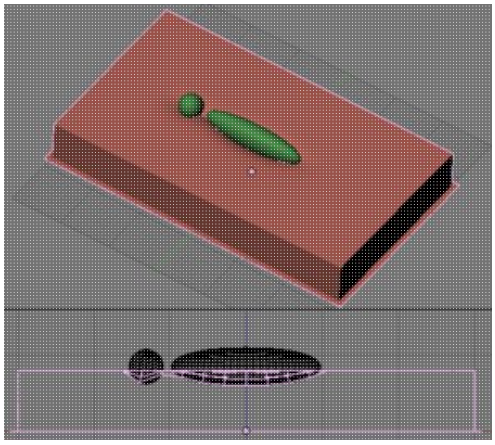


Image 3b: The result. Soft Body interactive modeling.

- Select the mattress.
- Apply the Soft Body modifier.

This removes the Soft Body modifier and makes the changes permanent.

- Now you can stop the animation playback!

You can use goals to either pin certain parts of a mesh to a location, or restrict the amount of movement away from the defined goal. So you may for example use armatures to animate your object but use a Goal group to define areas that are flexible and that react to Soft Body deformation. Since a simulation can be set up to take into account mass and gravity, you can easily create animated objects with realistic physical behavior using the Soft Body solver.

You can also use one of the following methods to combine Soft Bodies, objects and armatures:

- Soft Body vertices can be made the vertex parent of an empty, this empty can be the target of an IK solver.
- Using Bone stretching you can allow a bone to follow the length change of a Soft Body.
- You can use constraints to copy animated parameters of a bone onto an object.
- You can use a hook to bind a Soft Body to an armature driven object.

Swing

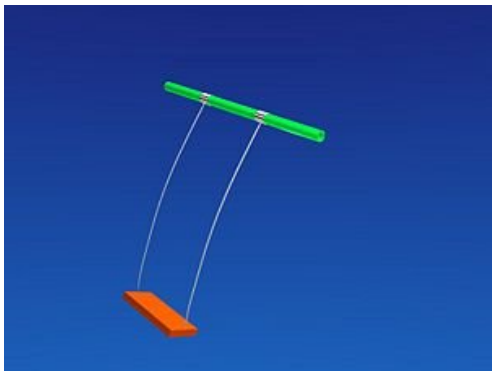


Image 1a: The actuator of the Swing is a Soft Body.
Animation
Blend file
The Soft Body simulation will drive the free end of a pendulum. This movement is used to drive the seat of a swing. You could also use a similar simulation for a Bungee cord, if you make the cord less stiff.

- Add a plane. Rename this object to â← `Pendulumâ←` .
- Change to Edit Mode.
- Make sure you are in top view and delete the left half of the plane, so that all that remains are two vertices connected by an edge.
- Select one vertex and make it the sole member of a vertex group (weight 1). This vertex will stay put. It is the â← upperâ← vertex.
- Bring the â← upperâ← vertex and the object center together, either by moving the vertices or the center.
- Tab out of Edit mode and enable Soft Body for the object.
- Use Goal, and select the vertex group of the â← upperâ← vertex.
- Set G Stiff to 0.

This creates a rig where the â← upperâ← vertex is pinned, but the other vertex (and therefore the edge) is fully affected by the simulation.

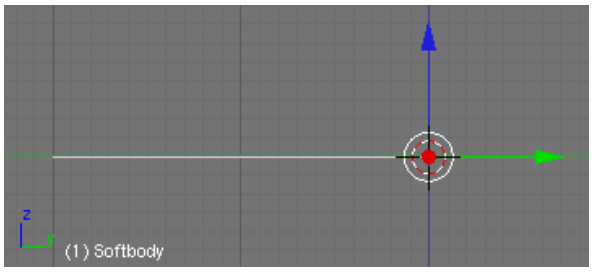


Image 1b: The Soft Body object.
When you scrub or play (AltA) your animation, the vertex associated with a goal stays put, but the other is free to be affected by the physics simulation. Since the edge is pretty stiff, and there is gravity and mass and thus inertia, the vertex at the other end swings back and forth, and will eventually come to a stop, even if you turn off all damping. If you donâ← t want it to stop, you could either increase the calculation accuracy, or use the RKCP solver, which is designed specifically for these types of simulations.

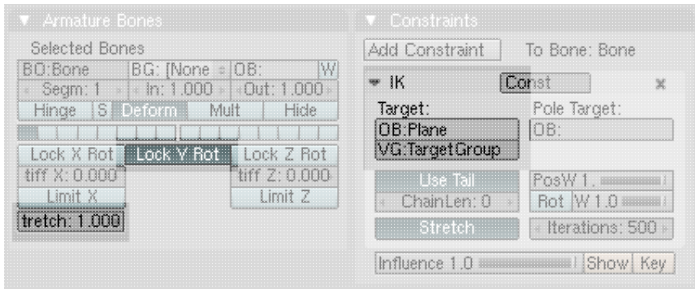


Image 1c: IK settings for the first bone.
The free vertex will become the IK target for an armature. Therefore we need to create a second vertex group.

- Change to Edit mode.
- Select the â← lowerâ← vertex.
- Press CtrlG â Add selected to new Group and name the vertex group â← `TargetGroup`â← .
- Change back to Object mode.
- Insert an Armature with a single bone.
- Ensure that the armature is selected and change to Edit mode.
- Move the tip of the bone to the â← lowerâ← vertex.
- Change to Pose mode and ensure that the bone is selected.
- Add an IK Solver to the bone (Editing context, Constraints panel, Add Constraint), use the object â← `Pendulum`â← and the vertex group â← `TargetGroup`â← as target.
- Set Stretch in the Armature Bones panel to 1.0.
- Lock Y Rot, so the bone doesnâ← t rotate around itself.

If you now play the animation, the bone swings back and forth. When the length of the object changes, the length of the bone also changes.

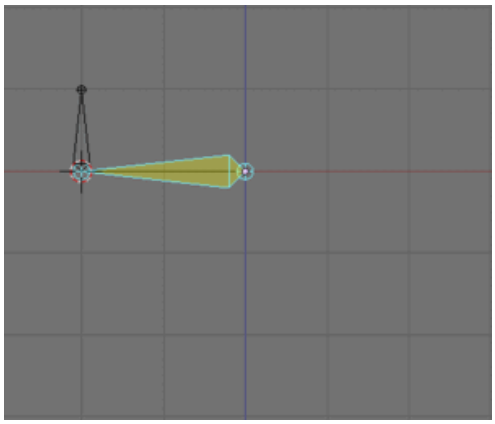


Image 1d: The armature is connected with the Soft Body by the IK Solver Target.

- Now extrude a second bone at an angle perpendicular to the first.
- Turn on S (Hinge Scale), so that the second bone isnâ← t scaled when the the first bone is.

Now you can parent an object, e.g. the seat of the swing, to this second bone.

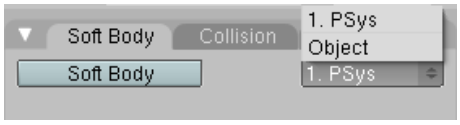


Image 1a: Activating Soft Body simulation for a Hair particle system.
Particle systems of type Hair may be animated as a softbody.

1. Create a working Hair particle system. It is not necessary to make it editable.
2. Change to the Physics sub-context.
3. Select the particle system instead of the object in the Soft Body panel (*Image 1a*).
4. Turn on Soft Body.

The default settings are not optimal for hair, so lets change them.



Image 1b: Soft Body hair with wind.
Animation
Blend file
Each hair has a number of keypoints (number of segments plus one) which perform the role of the vertices in Softbody calculation. It works in a similar way to animating curves with varying Control Point Weights, where the root gets a weight of 1.0, and the tip a weight of 0.0.

The *Control Point Weight* serves as the *Goal* for the Softbody simulation, so you can control how much each section of the hair moves. To edit the hair you need to set it Editable, then you can switch to Particle mode and perform weight painting on the control points.

To optimize the simulation you should:

1. Keep the number of hairs as low as possible.
2. Keep the number of segments as low as possible.
3. Avoid collisions as much as possible, e.g. work with a simplified object that calculates collisions.

Goal, Goal, Goal!

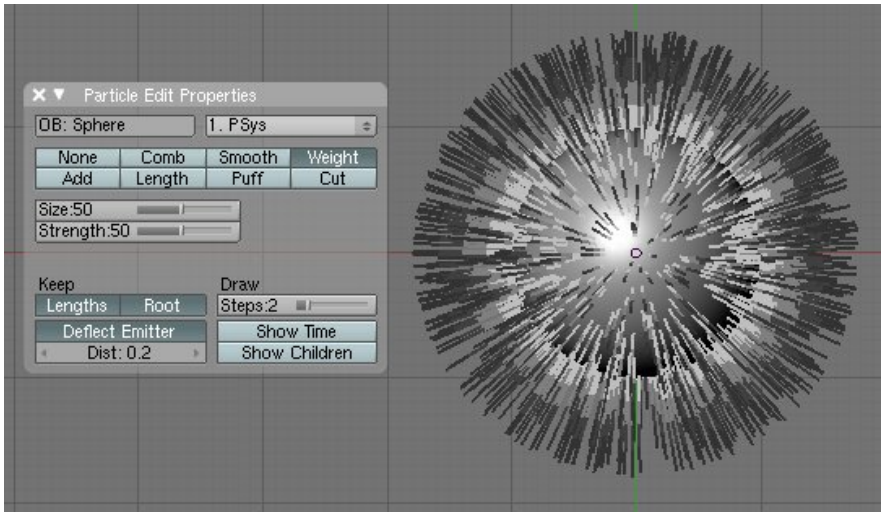


Image 2: Weight painting for Hair particles.
Each keypoint of a hair particle tries to maintain itâ← s position based on itâ← s weight setting, i.e. the hair tries to stay in shape. Imagine that there is a little spring which connects the keypoint with itâ← s original position. The goal stiffness (G Stiff) is the stiffness of this spring. The Goal value defines how much the keypoint responds to external forces (that have been configured to interact with the softbody animation).

- If the goal value of a keypoint is 1, the keypoint wonâ← t be affected by softbody animation (but of course is will still move if you apply normal animation to the object).
- If the goal value of a keypoint is 0, the keypoint is fully affected by softbody animation.

By default the root keypoints of hair particles have a goal of 1, so they will follow the emitter of the particles. The tip keypoints of hair particles have a goal of 0, so they are fully affected by the softbody animation. In between the root and tip weights are linearly interpolated.

- The root keypoint of a hair particle will always have a weight of 1, no matter what you do.
- You may paint the weight of each of the other keypoints individually in Particle mode (*Image 2*). A Strength of 50 is a weight of 0.5. So you can define parts of the hair that donâ← t move at all, e.g. behind the ears or held by a ribbon.
- The goal weight of each keypoint - except the root keypoint - is scaled by the G Min and G Max values (Physics sub-context, Soft Body panel). This is much quicker than trying to adjust the weights individually.

The force that keeps the keypoints together is calculated using the edge settings in softbody animation.

Turning jello into hair

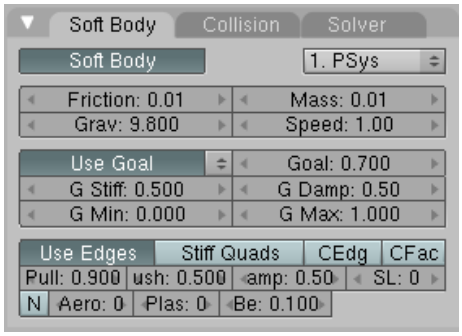


Image 3: No longer jello: these settings are a good starter for a softbody hair animation.

We are not doing an exact physical simulation, instead we are faking the look and behaviour of hair, but there are so many different hair types. So I will try and list some common properties (this is not exactly scientific).

1. Hair reacts to wind (if you only need wind you donâ← t necessarily need a softbody simulation, hair particles can react directly to force fields, like all other particles).
2. It follows the movement of the body.
3. It bends smoothly.
4. It doesnâ← t change in length.
5. It comes to rest quickly (doesnâ← t jitter).

So:

1. Hair needs some Friction (e.g. 0.01).
2. The mass should be lowered (e.g. to 0.01) unless you like that heavy silky hair advertisement look.
3. A little bending stiffness (Be to 0.1) helps the hair fall nicely.
4. Edge Pull should be quite high (as large as, or higher than 0.9).
5. It needs some serious damping (G Damp), e.g. 0.5 or higher.

If the mass of the keypoints is higher, you need more friction to stop the hair from wobbling.

G Min and G Stiff govern the appearance of the hair, often it is sufficient to change the stiffness. But that depends on the length of the hair. You will find that long hair needs more than 5 keypoints.

Collisions and Deflection

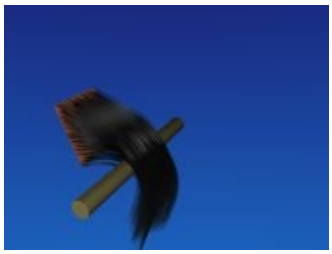


Image 4: Hair falling on a collision object.
Animation for *Image 4.* (info)
Let me give you the bad news first: no matter what you do, there is no way to achieve accurate deflection of hair when it collides with a fast moving object. If your object is accelerating quickly at least some parts of the hair will dive into it. The second problem is: collisions are slow to calculate. The third problem: collision is only calculated for the control points. So if you only have a few control points, there is very little data to base the collision on.

So avoid collision simulations as much as possible. Having said that, lets see what is possible with the interesting combination of hair and softbody.

- To let a softbody collide with another object, simply turn on Collision in the Collision panel of the Physics sub-context of the object.
- It often helps to increase the parameter Outer (the outer face thickness) of the collision object.
- When the keypoints of hair penetrate a collision object they are repelled. For hair this repulsion needs to be damped. If the repulsion is too strong increase Choke for the softbody system in the Solver panel.

As you can see, you can get a nice effect when hair falls slowly onto a collision object, though the deflection is not perfect (*Image 4*). If you do it the other way round and move the collision object slowly and at the

correct angle, the collision also works. If the collision object is moved too fast and/or at the wrong angle it simply moves through the hair.

Mode: Object mode

Panel: Object context â Physics sub-context â Soft Body, and Editing context â Modifiers

Hotkey: F7 to get to Object context; repeat to change sub-context.

A Soft Body simulation is started with AltA in the 3D window. The movement is then cached automatically (i.e. saved to disk), so you may have to clear the cache in the Collision panel if you have changed settings for collision objects (Free Cache or use CtrlB â Free Cache selected). You can also Bake the simulation, then you get a warning if you try to change a parameter which needs the cache to be cleared. You can also edit a baked simulation by hand.

Soft Body settings

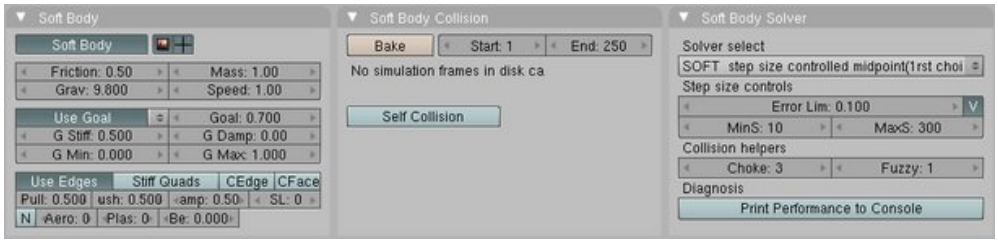


Image 1: The Soft Body panels.

General settings

Soft Body
: This enables the soft body modifier on the selected object.

Friction
: The friction of the surrounding medium. Generally friction dampens a movement.

Mass
: Mass value for vertices. Larger mass slows down acceleration, except for gravity where the motion is constant regardless of mass. Larger mass means larger inertia, so also braking a Soft Body is more difficult.

Grav
: Gravity, the amount of force in Z-axis direction. Earth gravity is a value close to 9.8 $m.s^{-2}$. Positive values make Softbodies drop; negative values make them rise.

Speed
: You can control the internal timing of the Softbody system with this value.

Start- and Endframe
The Start and End settings in the Collision panel are not only valid for the baking, but for all Soft Body simulations. So if your animation lasts longer than 250 frames, you have to increase the End value.

Goal settings

Use Goal
: Soft Body Goal acts like a pin on a chosen set of vertices; controlling how much of an effect soft body has on them. Enabling this tells Blender to use the position / animated position of a vertex in the simulation. Animating the vertices can be done in all the usual ways before the Soft Body simulation is applied. The *goal* is the desired end-position for vertices. How a softbody tries to achieve this goal can be defined using stiffness forces and damping.

Goal
: If no vertex group is used, this numeric field is the default goal weight for all vertices. If a vertex group is present and assigned, this button instead shows an popup selector button that allows you to choose the name of the goal vertex group. A Goal value of 1.0 means no Soft Body simulation, the vertex stays at its original (animated) position. When setting Goal to 0.0, the object is only influenced by physical laws. By setting goal values between 0.0 and 1.0, you can blend between having the object affected only by the animation system, and having the object affected only by the soft body effect.

G Stiff
: The spring stiffness for Goal. A low value creates very weak springs (more flexible â← attachmentâ← to the goal), a high value creates a strong spring (a stiffer â← attachmentâ← to the goal).

G Damp
: The friction for Goal. Higher values dampen the effect of the goal on the soft body.

G Min/G Max
: When you paint the values in vertex-groups (using Weight Paint mode), you can use the G Min and G Max to fine-tune (clamp) the weight values. The lowest vertex-weight (blue) will become G Min, the highest value (red) becomes G Max (please note that the blue-red color scale may be altered by User Preferences).

Edge settings

Use Edges
: The edges in a Mesh Object can act as springs as well, like threads in fabric.

 Pull
 : The spring stiffness for edges (how much the edges are stretched). A low value means very weak springs (a very elastic material), a high value is a strong spring (a stiffer material) that resists being pulled apart. 0.5 is latex, 0.9 is like a sweater, 0.999 is a highly-starched napkin or leather.

 Push
 : How much the Softbody resist being scrunched together, like a compression spring. Low values for fabric, high values for inflated objects and stiff material.

 Damp
 : The friction for edge springs. High values (max of 50) dampen the edge stiffness effect and calm down the cloth.

 SL (â← Shrink/Blowâ←)
 : The edges can shrink or been blown up. This value is given in percent, 0 disables this function. 100% means no change, the body keeps 100% of his size.

 Aero/N

If you turn on Aero the force is not confined to the vertices, but has an effect also on the edges. The angle and the relative speed between medium and edge is used to calculate the force on the edge. This force results that vertices with little connecting edges (front of a plane) fall faster than vertices with more connecting edges (middle of a plane). If all vertices have the same amount of edges in a direction they fall with equal speed. An edge moving in its own direction feels no force, and an edge moving perpendicular to its own direction feels maximum force (think of a straw moving through air). Try it with an Aero value of 30 at first. Enable the N button if you want to use an aerodynamic model that is closer to physical laws and looks more interesting. Disable for a more muted simulation.

Plas
: Plasticity, permanent deformation of the object.

Be (â← Bending Stiffnessâ←)
: This option creates virtual connections between a vertex and the one after the next. This includes diagonal edges. Damping applies also to these connections.

Stiff Quads
: For quad faces, the diagonal edges are used as springs. This stops quad faces to collapse completely on collisions (what they would do otherwise).

Sh (â← Shear stiffnessâ←)
: Stiffness of the virtual springs for quad faces.

CEdge
: Checks for edges of the softbody mesh colliding.

CFace
: Checks for any portion of the face of the softbody mesh colliding (compute intensive!). While CFace enabled is great, and solves lots of collision errors, there doesnâ← t seem to be any dampening settings for it, so parts of the softbody object near a collision mesh tend to â← jitterâ← as they bounce off and fall back, even when thereâ← s no motion of any meshes. Edge collision has dampening, so that can be controlled, but Deflection dampening value on a collision object doesnâ← t seem to affect the face collision.

Bake and Collision

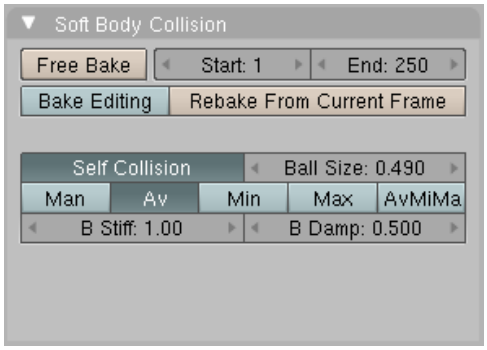


Image 2: Bake and Self Collision settings.

Bake

Bake
: Calculates the simulation and protects the cache. You need to be in Object mode to bake.

Free Bake
: Clears the cache.

Bake Editing
: You can edit the result of the simulation in Edit mode.

Rebake from Current Frame
: The manually edited result can be used as starting position for the next simulation run.

Start/End
: First and last frame of the simulation. Always valid, not only for *baking*.

Free Cache
: Clears the cache.

If you havenâ← t saved the blend file the cache is created in memory, so save your file first or the cache may be lost.

Self Collision

Self Collision is working only if you have activated Use Edges.

Self Collision
: When enabled, allows you to control how Blender will prevent the Soft Body from intersecting with itself. Every vertex is surrounded with an elastic virtual ball. Vertices may not penetrate the balls of other vertices. If you want a good result you may have to adjust the size of these balls. Normally it works pretty well with the default options.

Ball Size
: Default 0.49 BU or fraction of the length of attached edges. The edge length is computed based on the algorithm you choose. You know how when someone stands too close to you, and feel uncomfortable? We call that our â← personal spaceâ← , and this setting is the factor that is multiplied by the spring length. It is a spherical distance (radius) within which, if another vertex of the same mesh enters, the vertex starts to deflect in order to avoid a self-collision.
Set this value to the fractional distance between vertices that you want them to have their own â← spaceâ← . Too high of a value will include too many vertices all the time and slow down the calculation. Too low of a level will let other vertices get too close and thus possibly intersect because there wonâ← t be enough time to slow them down.
Man (â← manualâ←)
: The Ball Size directly sets the ball size (in BU).
Av (â← averageâ←)
: The average length of all edges attached to the vertex is calculated and then multiplied with the Ball Size setting. Works well with evenly distributed vertices.
Min/Max
: The ball size is as large as the smallest/largest spring length of the vertex multiplied with the Ball Size.
AvMiMax (â← average min/maxâ←)
: Size = ((Min + Max)/2) Ã Ball Size.

B Stiff
: Default 1.0. How elastic that ball of personal space is.

B Damp
: Default 0.5. How the vertex reacts. A low value just slows down the vertex as it gets too close. A high value repulses it.

Collisions with other objects are set in the (other) Collision panel. To collide with another object they have to share at least one common layer.

Soft Body Solver

These settings determine the accurateness of the simulation.

Blender has two solver algorithms that it can use:

Soft
: A step-size, adaptive algorithm for most situations.

RKCP (â← Runge Kutta Correct Physicsâ←)
: A mathematically and educationally correct algorithm. Good for something like a friction free pendulum, but tends to overshoot.

Error Limit/V
: Rules the overall quality of the solution delivered. Default 0.1. The most critical setting that says how precise the solver should check for collisions. Start with a value that is 1/2 the average edge length. If there are visible errors, jitter, or over-exaggerated responses, decrease the value. The solver keeps track of how â← badâ← it is doing and the Error Limit causes the solver to do some â← adaptive step sizingâ← . The V parameter helps the Solver figure out how much work it needs to do based on how fast things are moving.

MinS
: Minimum simulation steps per frame. Increase this value, if the Soft Body misses fast moving collision objects.

MaxS
: Maximum simulation steps per frame. Normally the number of simulation steps is set dynamically (with the Error Limit) but you have probably a good reason to change it.

Fuzzy
: Simulation is faster, but less accurate.

Choke
: Calms down (reduces the exit velocity of) a vertex or edge once it penetrates a collision mesh.

Print Performance to Console
: Prints on the console how the solver is doing.

Links

- Developer notes

If you are using a version of Blender up to 2.45 click here to access the old documentation.

Particles, Soft Bodies and Cloth objects can interact with their environment (other objects or particle systems). They may collide with them or be moved by forces (Fields). All types of objects and particles can generate fields, only mesh objects can be used as deflecting objects. Only curve object can bear Curve Guides.

- The objects need to share at least one common layer to have effect.
- You may limit the effect on particles to a group of objects (in the Extras panel).
- Fields can be generated by objects and/or by particles.

After changing the fields (Fields panel) or deflection (Collision panel) settings, you have to recalculate the particle, softbody or cloth system (Free Cache), this is not done automatically. You can clear the cache for all selected objects with CtrlB â Free cache selected.

Particles react to all kind of Force Fields, Soft Bodies only to Spherical/Wind/Vortex (they react on Harmonic fields but not in a useful way).

Mode: Object Mode

Panel: Object context â Physics sub-context â Fields

Hotkey: F7

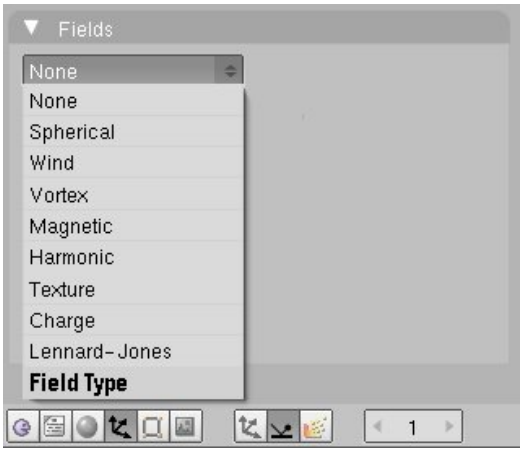


Image 1a: Fields panel in the Physics sub-context for an object without particles.

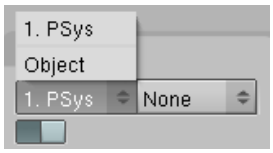


Image 1b: Fields panel for an object with particles.

To create a field you have to select the object and change to the Physics sub-context.

- Select the field type in the Fields panel. If the object has no particle system the panel looks like (*Image 1a*). An object can have only one field.
- If the object has one or more particle systems you can select one of the particle systems instead. Each particle system may have up to two fields assigned (*Image 1b*).

The fields have many options in common, these common options are explained for the Spherical field.

Spherical field

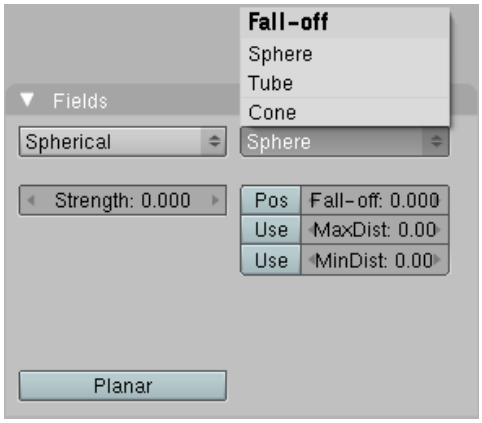


Image 2a: Force field panel for Spherical fields.
The Spherical field is the simplest of the fields. It gives a constant force towards (positive strength) or away from (negative strength) the objectâ← s center. Newtonian particles are attracted to a field with negative strength, and are blown away from a field with positive strength.

For Boids a field with positive strength can be used as a *Goal*, a field with negative strength can be used as *Predator*. Whether Boids seek or fly goals/predators depends on the Physics settings of the Boids.

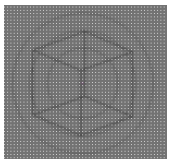


Image 2b: Spherical field indicator.

Strength
: The strength of the field effect. This can be positive or negative to change the direction that the force operates in. A force fieldâ← s strength is scaled with the force objectâ← s scale, allowing you to scale up and down scene, keeping the same effects. You can animate this parameter (Object Ipos, FStrength channel).

Planar
: The field is constant in the XY-plane, changes only in Z direction.

Fall-off
: Here you can specify the shape of the force field (if the Fall-off Power is not nil).
 Sphere
 : Spherical force field.
 Fall-off (Power)
 : How the power of the force field changes with the distance from the force field. If `r` is the distance from the center of the object, the force changes with $1/r^{Power}$. A Fall-off of 2 changes the force field with $1/r^2$. You can animate this parameter (Object Ipos, FFall channel).
 Pos
 : Fall-off only in the direction of the positive Z-Axis.
 Max Dist/Use
 : Makes the force field only take effect within a specified maximum radius (shown by an additional circle around the object). You can animate this parameter (Object Ipos, FMaxD channel).
 Min Dist/Use

The distance from the object center, up to where the force field is effective with full strength. If you have a Fall-off of 0 this parameter does nothing, because the field is effective with full strength up to Max Dist (or the infinity). Shown by an additional circle around the object.

Tube

Fall off results in a tube shaped force field. The force field can change *longitudinally* - Longitudinal, along the axis of the tube - or *radially* (Radial).

Cone

Fall off results in a cone shaped force field.

Other Field Types

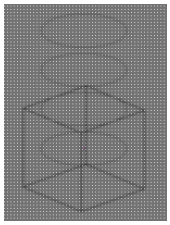


Image 3a: Wind field indicator.

Wind

Wind gives a constant force in a single direction, along the force objectâ← s local Z axis. The strength of the force is visualized by the spacing of the circles shown.

Noise

Random variation of the force. Nice, because you donâ← t have to use an IPO curve for that any more.

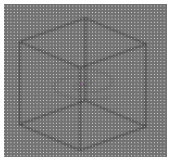


Image 3b: Vortex field indicator.

Vortex

Vortex fields give a spiraling force that twists the direction of points around the force objectâ← s local Z axis. This can be useful for making a swirling sink, or tornado, or kinks in particle hair.

Magnetic

This field depends on the speed of the particles.

Harmonic

The source of the force field is the zero point of a harmonic oscillator (spring, pendulum). If you set the Damping parameter to 1, the movement is stopped in the moment the object is reached. This force

field is really special if you assign it to particles. Normally every particle of the field system influences every particle of the target system. Not with Harmonic! Here every target particle is assigned to a field particle. So particles will move to the place of other particles, thus forming shapes. Tutorial: Particles forming Shapes.

Charge
: It is similar to spherical field except it changes behavior (attract/repulse) based on the effected particles charge field (negative/positive), like real particles with a charge. This mean this field has only effect on particles that have also a Charge field (else, they have no â← chargeâ← , and hence are unaffected)!

Lennard-Jones
: This field is a very short range force with a behavior determined by the sizes of the effector and effected particle. At a distance smaller than the combined sizes the field is very repulsive and after that distance itâ← s attractive. It tries to keep the particles at an equilibrium distance from each other. Particles need to be at a close proximity to each other to be effected by this field at all.

Particles can have for example both a charge and a Lennard-Jones potential - which is probably something for the nuclear physicists amongst us.

Texture field

You can use a texture force field to create an arbitrarily complicated force field, which force in the 3 directions is color coded. Red is coding for the x-axis, green for the y-axis and blue for the z-axis (like the color of the coordinate axes in the 3D window). A value of 0.5 means no force, a value larger than 0.5 acceleration in negative axis direction (like -Z), a value smaller than 0.5 acceleration in positive axis direction (like +Z).

Texture mode
: This sets the way a force vector is derived from the texture.

 RGB
 : Uses the color components directly as the force vector components in the color encoded directions. You need an RGB texture for this, e.g. an image or a colorband. So a Blend texture without a colorband would not suffice.

 Gradient
 : Calculates the force vector as the 3d-gradient of the intensity (grayscale) of the texture. The gradient vector always points to the direction of increasing brightness.

 Curl
 : Calculates the force vector from the curl of the 3d-rgb texture (rotation of rgb vectors). This also works only with a color texture. It can be used for example to create a nice looking turbulence force with a color clouds texture with perlin noise.

 Nabla
 : It is the offset used to calculate the partial derivatives needed for Gradient and Curl texture modes.

Use Object Coordinates
: Uses the emitter object coordinates (and rotation & scale) as the texture space the particles use. Allows for moving force fields, that have their coordinates bound to the location coordinates of an object.

Root TexCo
: The Root TexCo is useful for hair as it uses the texture force calculated for the particle root position for all parts of the hair strand.

2D
: The 2D button disregards the particles z-coordinate and only uses particles x&y as the texture coordinates.

Remember that only procedural texture are truly 3D.

Examples:

- A single colored texture 0.5/0.0/0.5 creates a force in the direction of the positive y-axis, e.g. hair is orientated to the y-axis.
- A blend texture with colorband can be used to created a force â← planeâ← . E.g. on the left side 0.5/0.5/0.5, on the right side 1.0/0.5/0.5 you have a force plane perpendicular to XY (i.e. parallel to Z). If you use an object for the coordinates, you can use the object to push particles around.
- An animated wood texture can be used to create a wave like motion.

Curve Guide field

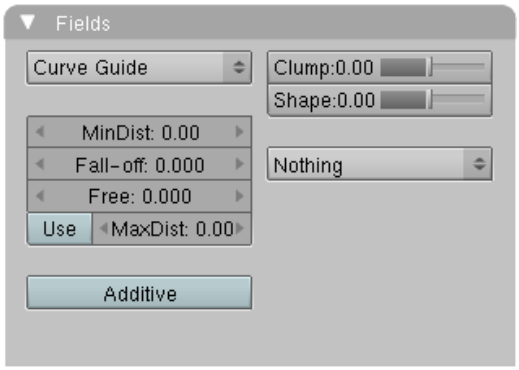


Image 4a: A Curve Guide field.
Curve objects can be the source of a Curve Guide field. You can guide particles along a certain path, they donâ← t affect Softbodys. A typical scenario would be to move a red blood cell inside a vein, or to animate the particle flow in a motor. You can use Curve Guides also to shape certain hair strands - though this may no longer be used as often now because we have the Particle Mode. Since you can animate curves as Softbody or any other usual way, you may build very complex animations while keeping great control and keeping the simulation time to a minimum.

The Curve needs the option CurvePath set to work. Blender sets this option when you assign a Curve Object this force field type automatically, but if you switch that off you have to restore it manually.

The option Curve Follow does not work for particles. Instead you have to set Angular Velocity (in the Physics panel of the Particle sub-context) to Spin and leave the rotation constant (i.e. donâ← t turn on Dynamic).

Curve Guides affect all particles on the same layer, independently from their distance to the curve. If you have several guides in a layer, their fields add up to each other (the way you may have learned it in your physics course). But you can limit their influence radius:

MaxDist/Use
: The maximum influence radius. Shown by an additional circle around the curve object.

MinDist

The distance from the curve, up to where the force field is effective with full strength. If you have a Fall-off of 0 this parameter does nothing, because the field is effective with full strength up to MaxDist (or the infinity). MinDist is shown with a circle at the endpoints of the curve in the 3D window.

Fall-off
: This setting governs the strength of the guide between MinDist and MaxDist. A Fall-off of 1 means a linear progression.

A particle follows a Curve Guide during itâ← s lifetime, the velocity depends from itâ← s lifetime and the length of the path. You can change that of course:

- A time Ipo curve of the emitter change the speed of the particles along the path, but time may not run backwards.

Free
: Fraction of particle life time, that is not used for the curve.

Additive
: If you use Additive, the speed of the particles is also evaluated depending on the Fall-off.

The other settings govern the form of the force field along the curve.

Clump
: The particles come together at the end of the curve (1) or they drift apart (-1).

Shape
: Defines the form in which the particles come together. +0.99: the particles meet at the end of the curve. 0: linear progression along the curve. -0.99: the particles meet at the beginning of the curve.

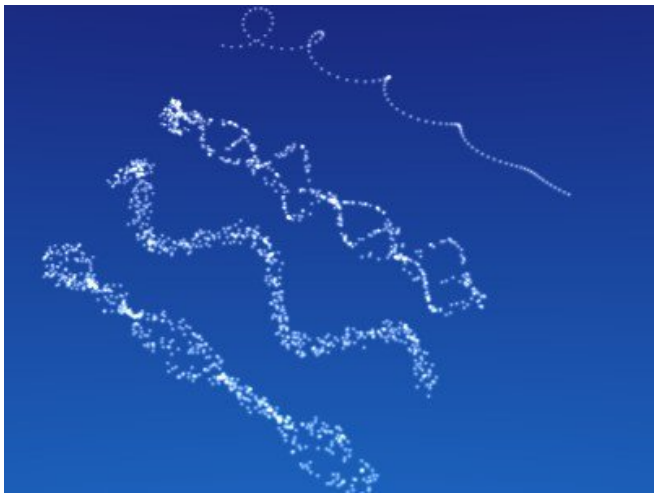


Image 4b: Kink options of a curve guide. From left to right: Radial, Wave, Braid, Roll.

Animation

With the drop down box Kink, you can vary the form of the force field:

Curl
: The radius of the influence depends on the distance of the curve to the emitter.

Radial
: A three dimensional, standing wave.

Wave

A two dimensional, standing wave.

Braid
: Braid.

Roll
: A one dimensional, standing wave.

It is not so easy to describe the resulting shapes, I hope itâ← s shown clearly enough in (*Image 4b*).

Links

- Wind & Deflector force update 2.48
- Particle options and guides (v2.40)

If you are using a version of Blender up to 2.45 click here to access the old documentation.

Particles, Soft Bodies and Cloth objects may collide with mesh objects. Boids try to avoid Collision objects.

- The objects need to share at least one common layer to have effect.
- You may limit the effect on particles to a group of objects (in the Extras panel).
- *Deflection* for softbody objects is difficult, they often penetrate the colliding objects.
- Hair particles ignore deflecting objects (but you can animate them as softbodies which take deflection into account).

If you change the deflection settings for an object you have to recalculate the particle, softbody or cloth system (Free Cache), this is not done automatically. You can clear the cache for all selected objects with CtrlB â Free cache selected.

Mode: Object Mode

Panel: Object context â Physics sub-context â Collision

Hotkey: F7

Options

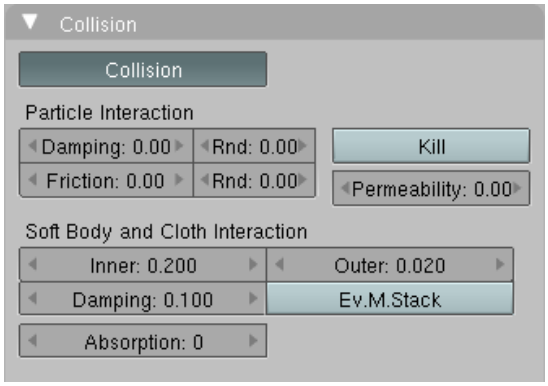


Image 1: Collision panel in the Physics sub-context.

Particle Interaction

Kill
: Ends the â← lifeâ← of colliding particles.

Damping
: Damping during a collision (independent of the velocity of the particles). Can be animated with Object Ipos, RDamp channel (*obviously, there is a mistake naming the curves*).
Rnd damping
: Random variation of damping. Can be animated with Object Ipos, Damping channel.

Friction
: Friction during movements along the surface.
Rnd friction
: Random variation of friction.

Permeability
: Fraction of particles passing through the mesh. Can be animated with Object Ipos, Perm channel.

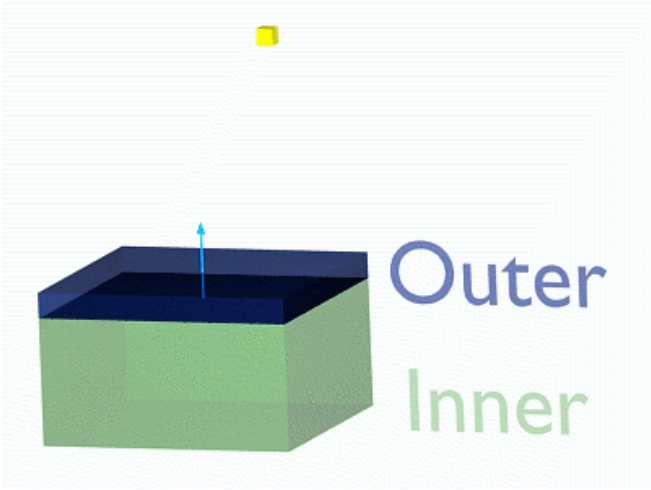


Image 1b: A softbody vertex colliding with a plane.

Soft Body and Cloth Interaction

Inner
: Size of the inner collision zone (padding distance).

Outer
: Size of the outer collision zone.

Outside and inside is defined by the face normal, depicted as blue arrow in (*Image 1b*).

Damping
: Damping during a collision.

Ev. M. Stack
: The modifier stack is evaluated first, then the collision is calculated. Enabling this is important if you want to take the SubSurf modifier into account (or donâ← t, because this may be really time

extensive).

Softbody collisions are difficult to get perfect. If one of the objects move too fast, the soft body will penetrate the mesh. See also the section about Soft Bodies.

Effector Interaction

Absorption
: A deflector can also deflect effectors. You can specify some collision/deflector objects which deflect a specific portion of the effector force using the Absorption value. 100% absorption results in no force getting through the collision/deflector object at all. If you have 3 collision object behind each other with e.g. 10%, 43% and 3%, the absorption ends up at around 50%
(`100Ã (1-0.1)Ã (1-0.43)Ã (1-0.03)`).

Examples



Image 2: Deflected Particles.
Here is a Meta object, dupliverted to a particle system emitting downwards, and deflected by a mesh cube:

Hints

- Make sure that the normals of the mesh surface are facing towards the particles/points for correct deflection.
- Hair particles react directly to force fields, so if you use a force field with a short range you donâ ← t need necessarily collision.
- Hair particles avoid their emitting mesh if you edit them in Particle mode. So you can at least model the hair with collision.

Starting with Blender 2.41 you can bake game engine hard body dynamics into animation Ipo curves.

Setting up world physics

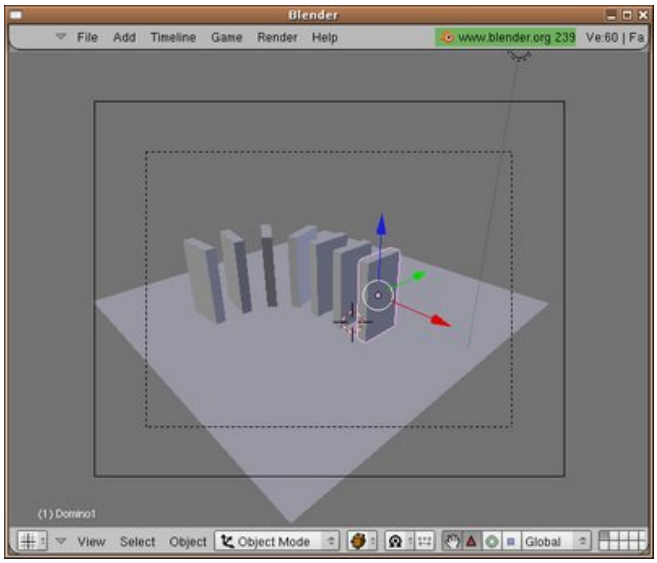
First set up the world physics. Go to Shading context (F5) [1], then the World sub-context [2], and control that

the physics engine is set to Bullet [3]:



Setting up the physics engine.

Setting up your scene



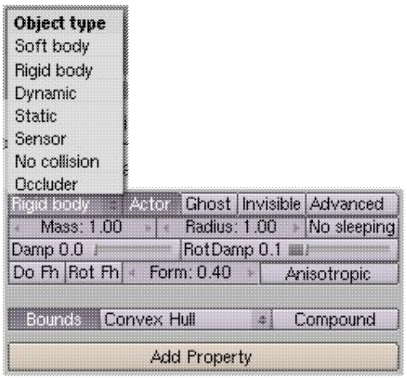
Domino scene.

Now set up your scene. In our case we will place a line of dominoes on a plane.

Create a plane (Add â Mesh â Plane) and some dominoes (Add â Mesh â Cube), arranged as shown.

Setting up collision objects

Next select one of the objects that will collide - in our case the first of the dominoes in the line. Tilt it toward the second (rotate) - weâ ll start it tipping over and let Blender do the rest.



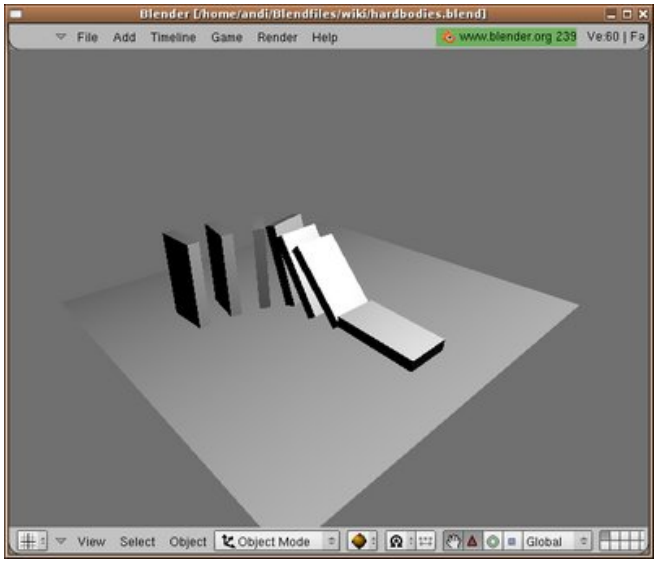
Dominoes physics settings.
Now for each of the objects (in this case, the dominoes) that will collide:

- Go to the Logic context (F4) and click on Actor.
- Select RigidBody in the drop down list.
- Click Bounds.
- Change the Collision Bounds type to Convex Hull (other types can be done but are more complex to setup).

Do this for each object that will be involved in the collisions. The plane (ground) in this case should not have its logic Actor and Bounds set. Although it does interact with the other objects (in this case, supporting the dominoes), if we set it as a dynamic it too would fall away by gravity. So, it will be a static actor.

Also do not forget that physics objects **have** to have a material, or they will bounce around wildly because they need the damping values etc. which are not set on objects without materials.

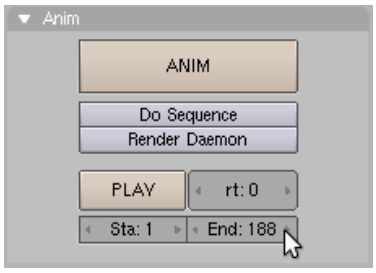
Animation preview



Previewing the animation with P.
Press P to see a preview of the game animation.

Press Esc when done.

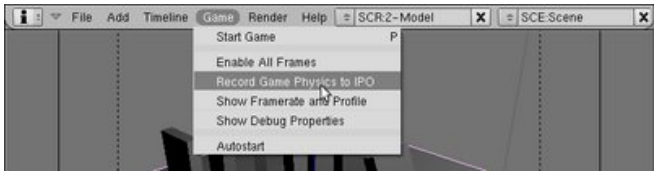
Get ready to bake



Anim panel.
Go to the Scene context (F10).

Select the Keyframe start and end range you want the animation you record to start and end on.

Record the Ipo

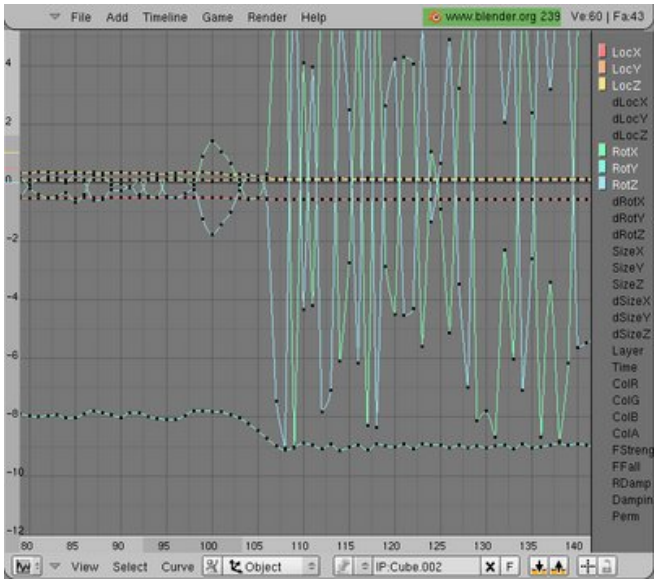


Recording game physics to Ipos.
Select Game â Record Game Physics to IPO.

Press the P.

Press Esc when done

Scrub through the animation



Ipos of the baked game physics.

Now you can scrub forward and backwards through your animation with the arrow keys, and look at the animation curves in Ipo Curve Editor by selecting an object.

Replace the animation

If you change something and press the P the old animation will be recorded over if you have Record Game Physics to IPO still enabled.

Troubleshooting and hints

As with all physics systems there are a number of limitations that you need to be aware of - please see this page for tips, tricks, and limitations.

Fluid Simulation

Mode: Object mode / Edit mode (Mesh)

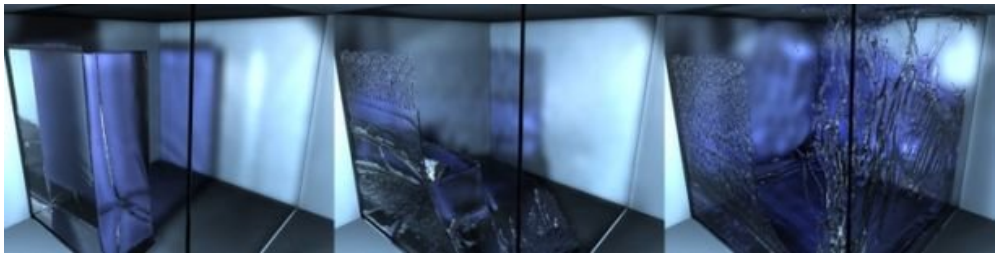
Panel: Physics sub-context â Fluid

Hotkey: F7

Description

While modeling a scene with blender, certain objects can be marked to participate in the fluid simulation, e.g. as fluid or as an obstacle. The bounding box of another object will be used to define a box-shaped region to simulate the fluid in (the so called â← simulation domainâ←). The global simulation parameters (such as viscosity and gravity) can be set for this domain object.

Using the BAKE button, the geometry and settings are exported to the simulator and the fluid simulation is performed, generating a surface mesh together with a preview for each animation frame, and saving them to hard disk. Then the appropriate fluid surface for the current frame is loaded from disk and displayed or rendered.



A breaking dam example.

Process

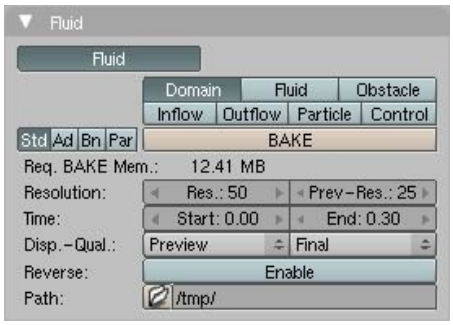
In general, you follow these steps:

1. Model the scene (objects, materials, lights, camera).
2. Designate the portion of the scene where the fluid will flow (the domain).
3. Specify the functions of the various objects as they relate to the fluid (an inlet, outlet, or obstruction, etc.).
4. Create the fluid source(s), and specify its (their) material, viscosity, and initial velocity.
5. Bake in a preliminary simulation.
6. Revise as necessary, saving changes.
7. Bake in a final simulation.

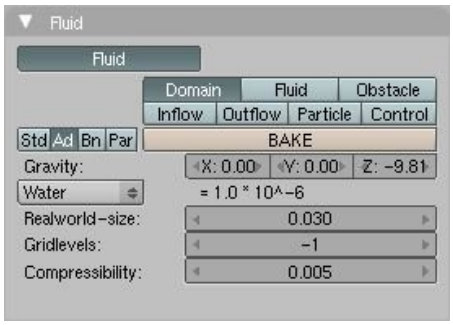
When you Bake, the version identifier in the User Preferences header changes to a status bar, showing the number of frames baked and the total number to go. Baking takes a **lot** of compute power (and time). Donâ← t watch it, because we all know that a watched pot does not boil, a watched cake does not bake, and watching fluid bake is like watching grass grow. Baking is best done overnight.

Because of the possibility of spanning and linking between scenes, there can only be one domain in an entire .blend file.

Options



The basic (and frequently needed) fluid simulation options.

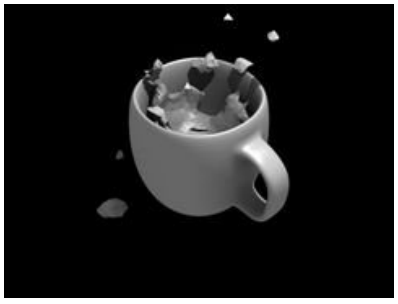


Less frequently needed advanced options.

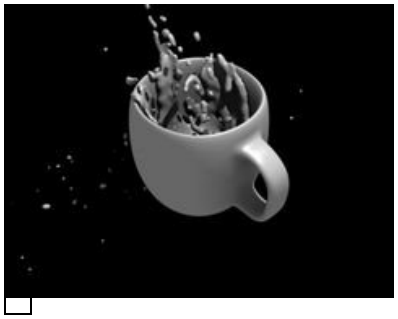
Domain/Fluid/Obstacle/Inflow/Outflow/Particle/Control
: Selecting one of these buttons determines how the enabled object will be used during the simulation. Each button determines a different functionality/interaction, and each has different further options that will become available.

Mesh rendering
If the mesh has modifiers, the rendering settings are used for exporting the mesh to the fluid solver. Depending on the setting, calculation times and memory use might exponentially increase. For example, when using a moving mesh with Subsurf as an obstacle, it might help to decrease simulation time by switching it off, or to a low subdivision level. When the setup/rig is correct, you can always increase settings to yield a more realistic result.

Domain



10cm mug at Resolution 70.



10cm mug at Resolution 200.
The bounding box of the object serves as the boundary of the simulation. **All fluid objects must be in the domain.** Fluid objects outside the domain will not bake. No tiny droplets can move outside this domain; itâ← s as if the fluid is contained within the 3D space by invisible force fields. There can be only a single fluid simulation domain object in the scene. The lengths of the bounding box sides can be different.

The shape of the object does not matter because it will *always* be a fixed cubic of space. So, usually there wonâ← t be any reason to use another shape than a box. If you need obstacles or other boundaries than a box to interfere with the fluid flow, you need to insert additional obstacle objects *inside* the domain boundary.

This object will be *replaced* by the fluid during the simulation.

Resolution

The granularity at which the actual fluid simulation is performed. This is probably the most important setting for the simulation as it determines the amount of details in the fluid, the memory and disk usage as well as computational time.

Note that the amount of required memory quickly increases: a resolution of 32 requires ca. 4MB, 64 requires ca. 30MB, while 128 already needs more than 230MB. Make sure to set the resolution low enough, depending on how much memory you have, to prevent Blender from crashing or freezing. Remember also that many operating systems limit the amount of memory that can be allocated by a single *process*, such as Blender, even if the *machine* contains much more than this. Find out what limitations apply to your machine.

Be sure to set the resolution appropriate to the real-world size of the domain (the actual physical size of the domain, in meters, is set in the Advanced Settings). If the domain is not cubic, the resolution will be taken for the longest side. The resolutions along the other sides will be reduced according to their lengths (therefore, a non-cubic domain will need less memory than a cubic one, resolutions being the same).

Preview-Res

This is the resolution at which the preview surface meshes will be generated. So it does not influence the actual simulation. Even if â← there is nothing to seeâ← in the preview, there might be a thin fluid surface that cannot be resolved in the preview.

Start and End time

Start is the simulation start time (in seconds). This option makes the simulation computation in Blender start later in the simulation. The domain deformations and fluid flow prior to the start time are not saved.

For example, if you wanted the fluid to appear to already have been flowing for 4 seconds before the actual first frame of data, you would enter 4.0 here.

End is the simulation ending time (in seconds).

Start and end times have nothing to do with how many frames are baked, but instead are based on physical force and fluid viscosity.

Baking always starts at Frame #1:
The fluid simulator disregards the Sta setting in the Anim panel, it will always bake from frame 1. If you wish the simulation to start later than frame 1, you must key the fluid objects in your domain to be inactive until the frame you desire to start the simulation. See below for more information.

Baking always ends at the End Frame set in the Anim panel:
If your frame-rate is 25 frames per second, and ending time is 4.0 seconds, then you should (if your start time is 0) set your animation to end at frame `4.0 Ã  25 =` 100.

â← Start Timeâ← and â← End Timeâ← are expressed in *simulation* time (seconds):
If you set Start time to 3.0, and End time to 4.0, you will simulate 1 second of fluid motion. That one second of fluid motion will be spread across however-many frames are set in your End setting in the Anim panel.
This means, for example, that if you have Blender set to make 250 frames at 25 fps (Scene context â Render sub-context â Anim and Output panels), the fluid will look like it had already been flowing for 3 seconds at the start of the simulation, *but* will play in slow motion (one-tenth normal speed), since the 1 second fluid sim plays out over the course of 10 video seconds. To correct this, change the end time to `3.0 + 10.0 =` 13.0 to match the 250 frames at 25 fps. Now, the simulation will be real-time, since you set 10 seconds of fluid motion to simulate over 10 seconds of animation. Having these controls in effect gives you a â← speed controlâ← over the simulation.

Disp.-Qual

How to display a baked simulation in the Blender GUI (first pulldown menu) and for rendering (second one):

- Geometry : Use the original geometry (before simulation).
- Preview : Use the preview mesh.
- Final : Use the final high definition mesh.

When no baked data is found, the original mesh will be displayed by default.

After you have baked a domain, it is displayed (usually) in the Blender window as the preview mesh. To see the size and scope of the original domain box, select Geometry in the left dropdown.

Bake directory

REQUIRED! Directory and file prefix to store baked surface meshes with. This is similar to the animation output settings, only selecting a file is a bit special: when you select any of the previously generated surface meshes (e.g. â← `untitled_OBcube_fluidsurface_final_0132.bobj.gz`â←), the prefix will be automatically set (â← `untitled_OBcube_`â← for this example). This way the simulation can be done several times with different settings, and allows quick changes between the different sets of surface data.

The default value is â← `/tmp/`â← , which is probably *not* what you want. Choose an appropriate directory-name and file prefix so that these files will be stored in an appropriate location *and* named in such a way that two different fluid-simulations wonâ← t conflict with one another (if youâ← re intending to specify only a directory-name here, i.e. without a filename-prefix, donâ← t forget the trailing â← `/`â←).

â← BAKEâ← Button

Perform the actual fluid simulation. The Blender GUI will freeze and only display the current frame that is simulated. Pressing Esc will abort the simulation. Afterwards two â← `.bobj.gz`â← (one for the Final quality, one for the Preview quality), plus one â← `.bvel.gz`â← (for the Final quality) will be in the selected output directory for each frame.

Notesâ←¦

Freeing the previous baked solutions

Deleting the content of the â← Bakeâ← directory is a destructive way to achieve this. Be careful if more than one simulation uses the same bake directory (be sure they use different filenames, or they will overwrite one another)!

Reusing Bakes

Manually entering (or searching for) a previously saved (baked) computational directory and filename mask will switch the fluid flow and mesh deformation to use that which existed during the old bake. Thus, you can re-use baked flows by simply pointing to them in this field.

Selecting a Baked Domain

After a domain has been baked, it changes to the fluid mesh. To re-select the domain so that you can bake it again after you have made changes, go to any frame and select (RMB) the fluid mesh. Then you can click the BAKE button again to recompute the fluid flows inside that domain.

â← Stâ← /â← Adâ← /â← Bnâ← /â← Parâ← Buttons

Till now, we were in the Standard buttons. Clicking another one of these buttons will show other â← panelsâ← (groups of controls: Advanced, Bn for boundary, and Particle) of more advanced options, that often are fine set at the defaults.

Standard

The settings in this set are already been described aboveâ←¦

Advanced

Gravity vector

Strength and direction of the gravity acceleration and any lateral (x,y plane) force. The main component should be along the negative z-axis (in $m.s^{-2}$).
Please Note: All of the x,y,z values should not be zero, or the fluid wonâ← t flow! Imagine a droplet floating in the nothingness of deep spaceâ←¦ It must be some small number in at least one direction.

Viscosity

The â← thicknessâ← of the fluid and actually the force needed to move an object of a certain surface area through it at a certain speed. You can either enter a value directly or use one of

the presets in the drop down (such as honey, oil, or water).
For manual entry, please note that the normal real-world viscosity (the so-called dynamic viscosity) is measured in Pascal-seconds (Pa.s), or in Poise units (P, equal to 0.1 Pa.s, pronounced â← *pwaz*â← , from the Frenchman Jean-Louis Poiseuille, who discovered the laws on â← the laminar flow of viscous fluidsâ←), and commonly centiPoise units (cP, equal to 0.001 Pa.s, â← *sentipwaz*â←). Blender, on the other hand, uses the kinematic viscosity (which is dynamic viscosity in Pa.s, divided by the density in kg.m^{-3}, unit $m^2.s^{-1}$). The table below gives some examples of fluids together with their dynamic and kinematic viscosities. Manual entries are specified by a floating point number and an exponent. These floating point and exponent entry fields (scientific notation) simplify entering very small or large numbers. The viscosity of water at room temperature is 1.002 cP, ou 0.001002 Pa.s; the density of water is about 1000 kg.m^{-3}, which gives us a kinematic viscosity of 0.000001002 m^2.s^{-1} - so the entry would be 1.002 times 10 to the minus six ($1.002Ã\ 10^{-6}$ in scientific notation). Hot Glass and melting iron is a fluid, but very thick; you should enter something like $1.0Ã\ 10^0$ (= 1.0) as its kinematic viscosity (indicating a value of $1.0Ã\ 10^6$ cP).
Note that the simulator is not suitable for non-fluids, such as materials that do not â← flowâ← . Simply setting the viscosity to very large values will not result in rigid body behavior, but might cause instabilities.

Viscosity varies
The default values in Blender are considered typical for those types of fluids and â← look rightâ← when animated. However, actual viscosity of some fluids, especially sugar-laden fluids like chocolate syrup and honey, depend highly on temperature and concentration. Oil viscosity varies by SAE rating. Glass at room temperature is basically a solid, but glass at 1500 degrees Celsius flows (nearly) like water.

Blender Viscosity Unit Conversion

Fluid	dynamic viscosity (in cP)	kinematic viscosity (Blender, in m^2.s^{-1})
Water (20Â°C)	1.002Ã 10 ^{0} (1.002)	1.002Ã 10 $^{-6}$ (0.000001002)
Oil SAE 50	5.0Ã 10 ^{2} (500)	5.0Ã 10 $^{-5}$ (0.00005)
Honey (20Â°C)	1.0Ã 10 ^{4} (10,000)	2.0Ã 10 $^{-3}$ (0.002)
Chocolate Syrup	3.0Ã 10 ^{4} (30,000)	3.0Ã 10 $^{-3}$ (0.003)
Ketchup	1.0Ã 10 ^{5} (100,000)	1.0Ã 10 $^{-1}$ (0.1)
Melting Glass	1.0Ã 10 ^{15}	1.0Ã 10 ^{0} (1.0)

Realworld-size
Size of the domain object in the real world in meters. If you want to create a mug of coffee, this might be 10 cm (0.1 meters), while a swimming pool might be 10m. The size set here is for the longest side of the domain bounding box.

Gridlevel
How many adaptive grid levels to be used during simulation - setting this to -1 will perform automatic selection.

Compressibility
If you have problems with large standing fluid regions at high resolution, it might help to reduce this number (note that this will increase computation times).

Bn (Boundaries)

Boundary type

This is the same as for obstacle objects below, it will basically set the six sides of the domain to be either sticky, non-sticky, or somewhere in between (this is set by the PartSlip Amount value).

Surface Smoothing

Amount of smoothing to be applied to the fluid surface. 1.0 is standard, 0 is off, while larger values increase the amount of smoothing.

Surface Subdivision

Allows the creation of high-res surface meshes directly during the simulation (as opposed to doing it afterwards like a subdivision modifier). A value of 1 means no subdivision, and each increase results in one further subdivision of each fluid voxel. The resulting meshes thus quickly become large, and can require large amounts of disk space. Be careful in combination with large smoothing values - this can lead to long computation times due to the surface mesh generation.

Generate SpeedVecs/Disable

If this button is clicked, no speed vectors will be exported. So by default, speed vectors are generated and stored on disk. They can be used to compute image based motion blur with the compositing nodes.

Particles

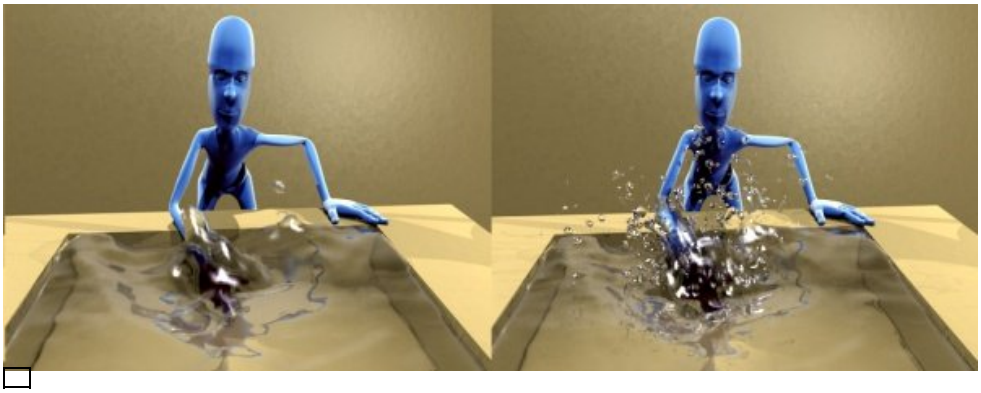
Here you can add particles to the fluid simulated, to enhance the visual effect.

Tracer Particles

Number of tracer particles to be put into the fluid at the beginning of the simulation. To display them create another object with the Particle fluid type, explained below, that uses the same bake directory as the domain.

Generate Particles

Controls the amount of fluid particles to create (0=off, 1=normal, >1=more). To use it, you have to have a surface subdivision value of at least 2.



An example of the effect of particles can be seen here - the image to the left was simulated without, and the right one with particles and subdivision enabled.

Fluid

All regions of this object that are inside the domain bounding box will be used as actual fluid in the simulation. If you place more than one fluid object inside the domain, they should currently not intersect. Also make sure the surface normals are pointing outwards. In contrast to domain objects, the actual mesh geometry

is used for fluid objects.

Volume init type

- Volume will initialize the inner part of the object as fluid. This works only for closed objects.
- Shell will initialize only a thin layer for all faces of the mesh. This also works for non closed meshes.
- Both combines volume and shell - the mesh should also be closed. See the picture below.



Example of the different volume init types: Volume, Shell and Both. Note that the shell is usually slightly larger than the inner volume.

Initial velocity
: Speed of the fluid at the beginning of the simulation, in meters per second.



The direction of Surface Normals makes a big difference!:

Blender uses the orientation of the Surface Normals to determine what is â← inside ofâ← the Fluid object and what is â← outsideâ← . You want all of the normals to face *outside* (in Edit mode, use CtrlN or press Space and choose Edit â Normals â Calculate Outside). If the normals face the wrong way, youâ← ll be rewarded with a â← gigantic flood of waterâ← because Blender will think that the volume of the object is outside of its mesh! This applies regardless of the Volume init type setting.

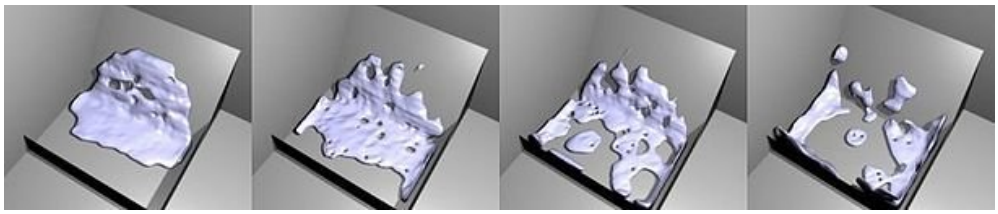
Obstacle

This object will be used as an obstacle in the simulation. As with a fluid object, obstacle objects currently should not intersect. As for fluid objects, the actual mesh geometry is used for obstacles. For objects with a volume, make sure that the normals of the obstacle are calculated correctly, and radiating properly (use the Flip Normal button, in Edit mode, Mesh Tools panel, Editing context [F9]), particularly when using a spinned container. Applying the Modifier SubSurf before baking the simulation could also be a good idea if the mesh is not animated.

Volume init type
: Same as for a fluid object above.

Boundary type (see picture below)
: Determines the stickiness of the obstacle surface, called â← Surface Adhesionâ← . Surface Adhesion depends in real-world on the fluid and the graininess or friction/adhesion/absorption qualities of the surface.
 - ◊ Noslip causes the fluid to stick to the obstacle (zero velocity).
 - ◊ Free(-slip) allows movement along the obstacle (only zero normal velocity).
 - ◊ Part(-slip) mixes both types, with 0 being mostly noslip, and 1 being identical to freeslip.

Note that if the mesh is moving, it will be treated as noslip automatically.



Example of the different boundary types for a drop falling onto the slanted wall. From left to right: no-slip, part-slip 0.3, part-slip 0.7 and free-slip.

Animated Mesh/Export
: Click this button if the mesh is animated (e.g. deformed by an armature, shape keys or a lattice). Note that this can be significantly slower, and is not required if the mesh is animated with position or rotation Ipos (i.e. only *object* transformations).

PartSlip Amount
: Amount of mixing between no- and free-slip, described above.

Animated Objects are No Slip:

If an obstacle is moving, Blender treats it automatically as no slip (sticky). If you want fluid to splash off of a moving object, put an transparent plane in the spot where the fluid will hit the moving object, exactly aligned and shaped as the object, but just a tiny bit in front of the object between the object and the fluid. As the object whizzes by and the fluid splashes, the splash will actually contact the stationary transparent plane and slide off, and the object will continue on its merry way.

Impact Factor
: Amount of fluid volume correction for gain/loss from impacting with moving objects. If this object is not moving, this setting has no effect. However, it if is and the fluid collides with it, a negative value takes volume away from the Domain, and a positive number adds to it. Ranges from -2.0 to 10.0.

Inflow

This object will put fluid into the simulation (think of a water tap).

Volume init type
: Same as for a fluid object above.

Inflow velocity
: Speed of the fluid that is created inside of the object.

Local Coords/Enable
: Use local coordinates for the inflow. This is useful if the inflow object is moving or rotating, as the inflow stream will follow/copy that motion. If disabled, the inflow location and direction do not change.

Animated Mesh/Export
: Same as for an obstacle object above.

Outflow

Any fluid that enters the region of this object will be deleted (think of a drain or a black hole). This can be useful in combination with an inflow to prevent the whole domain from filling up. When enabled, this is like a tornado (waterspout) or â← wet vacâ← vacuum cleaner, and the part where the fluid disappears will follow the object as it moves around.

Volume init type
: Same as for a fluid object above.

Animated Mesh/Export
: Same as for an obstacle object above.

Particle

This type can be used to display particles created during the simulation. For now only tracers swimming along with the fluid are supported. Note that the object can have any shape, position or type - once the particle button is pressed, a particle system with the fluid simulation particles will be created for it at the correct position. When moving the original object, it might be necessary to delete the particle system, disable the fluidsim particles, and enable them again. The fluidsim particles are currently also unaffected by any other particle forces or settings.

Particle type

Drops
: Surface splashes of the fluid result in droplets being strewn about, like fresh water, with low Surface Tension.

Floats
: The surface tension of the fluid is higher and the fluid heavier, like cold seawater and soup. Breakaways are clumpier and fall back to the surface faster than Drops, as with high Surface Tension.

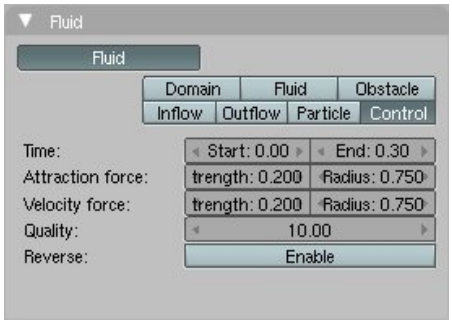
Tracer
: Droplets follow the surface of the water where it existed, like a fog suspended above previous fluid levels. Use this to see where the fluid level has been.

Size Influence
: The particles can have different sizes, if this value is 0 all are forced to be the same size.

Alpha Influence
: If this value is >0, the alpha values of the particles are changed according to their size.

Path (bake directory)
: The simulation run from which to load the particles. This should usually have the same value as the fluid domain object (e.g. copy by CtrlC, CtrlV).

Control



Fluid control options.

Time
: You specify the start and end time during which time the fluid control object is active.

Attraction force
: The attraction force specifies the force which gets emitted by the fluid control object. Positive force results in attraction of the fluid, negative force in avoidance.

Velocity force
: If the fluid control object moves, the resulting velocity can also introduce a force to the fluid.

Quality
: Higher quality result in more control particles for the fluid control object.

Reverse
: The control particle movement gets reversed.

See these release notes for more information.



Another example animation of a falling drop, simulated in Blender and rendered in Yafray.

Animating Fluid Property Changes

A new type of Ipo Curve, FluidSim, is available for fluid domain objects. Unlike most other animatable values in Blender, FluidSim Ipos cannot be keyframed by simply using the I key; you must manually set values by clicking in the Ipo window. In order to set a keyframe, you must select the property you wish to animate in the Ipo window and CtrlLMB click to set the keyframe to the desired location in the Ipo window.



Enter Properties:

Note that you do not have to be exact on where you click; we recommend that after you set the control point, open the Transform Properties panel (N) and round the X value to a whole frame number, and then set the Y value that you wish.

The fluid domain has several channels that control the fluid over time:

Fac-Visc
: A multiplicative factor in the fluidâ← s viscosity coefficient. It must be set before baking, and changes the viscosity of the fluid over time, so you can turn water into wiâ←¦ oil, for example!

Fac-Tim
: Changes the speed of the simulation; like the Speed Control in the VSE can speed up or slow down a video, this curve can speed up or slow down the fluid motion during a frame sequence. If the value for Fac-Tim is less than or equal to zero, time (and the fluid) stands still; the fluid freezes. For values between 0.0 and 1.0, the fluid runs slower and will look thicker. 1.0 is normal fluid motion, and values greater than 1.0 will make the fluid flow faster, possibly appearing thinner.

GravX/GravY/GravZ
: The XYZ vector for gravity changes; aka inertia of the fluid itself (think drinking a cup of coffee while driving NASCAR at Talladega, or sipping an espresso on the autobahn, or watering the plants on the Space Shuttle). Changes in these curves make the fluid slosh around due to external forces.

The Fluid, Obstacle, Inflow, Outflow and Particle objects can use the following channels:

VelX/VelY/VelZ
: Spurts of water from the garden hose can be simulated via these curves, to mimic changes in pressure and/or direction. It also can be used to simulate the effect of wind on a stream of water, for example.

Active
: When Active transitions from 0.0 to something greater than 0 (such as between 0.1 and 1.0), the objectâ← s function (designated as an Inflow, or Outflow, etc.) resumes its effect. Crossing down to 0.0 and then at some point, back up, re-establishes the effect and the resulting fluid sim. Use this for dripping, or any kind of intermittent inflow. This active status also works for objects designated as Outflow and Obstacle, so you can also simulate (for example) a drain plugging up.

You can also control the force settings of Control objects:

AttrForceStr, AttrForceRa
: These curves control the values of the attraction force settings.

VelForceStr, VelForceRa
: These curves control the values of the velocity force settings.

Technical Details



â← My cup runneth overâ← , created with Blender and Yafray.
Fluid animation can take a lot of time - the better you understand how it works, the easier it will be to estimate how the results will look. The algorithm used for Blenderâ← s fluid simulation is the *Lattice Boltzmann Method* (LBM); other fluid algorithms include *Navier-Stokes* (NS) solvers and *Smoothed Particle Hydrodynamics* (SPH) methods. LBM lies somewhere between these two.

In general, it is really hard for current computers to correctly simulate even a 1-meter tank of water. For simulating a wave crashing through a city, you would probably need one of the most expensive supercomputers you could get, and it might still not work properly, no matter which of the three algorithms above youâ← re using. Therefore, to achieve â← the effect that you really wantâ← , youâ← ll need to resort to strategies very similar to what filmmakers have been doing (quite successfullyâ←¦) in â← analogueâ← movies for many years: â← *fake it!*â←

A good fluid simulation is a *very important* part, but not the *only* part, of achieving a satisfactory image. Let Blender do the computational dirty-work of calculating the basic fluid simulation, then create realism by adding carefully selected details that match the viewerâ← s expectations for â← the real-life maelstrom that you have createdâ← .

For example, you can pretend to have a wave in a gigantic city by: building a *smaller* model, modeling a *small* wave in the model at very high resolution, and hope that nobody will notice the difference between a 100m and a 1m wave (they wonâ← t). Texture the wave front with lots of noise and clouds affecting the

color. Add lots of smoke (mist) emitters on the various surfaces that the wave hits, timing each of them to emit at the moment of impact in a direction incident to the surface and collision. Animate cars and trash (and drowning peopleâ←¦) to float and bob on the wave front using the baked mesh. Use a string of mist emitters pointing up positioned at the wave crest to simulate the mist that blows off the top of the crest into the air. Consider exactly where you want to put the camera, whether you want to use a zoom lens or a wide angle, and so on (is the viewer to be â← looking down upon the poor unfortunate actorsâ← , or â← drowning along with themâ← ?). *This* is the kind of attention to detail, above and beyond the fluid simulation itself, that will carry the shot.

For Blenderâ← s LBM solver, the following things will make the simulation harder to compute:

- Large domains.
- Long duration.
- Low viscosities.
- High velocities.

The viscosity of water is already really low, so especially for small resolutions, the turbulence of water can not be correctly captured. If you look closely, most simulations of fluids in computer graphics do not yet look like real water as of now. Generally, donâ← t rely on the physical settings too much (such as physical domain size or length of the animation in seconds). Rather try to get the overall motion right with a low resolution, and then increase the resolution as much as possible or desired.

Hints

- Donâ← t be surprised, but youâ← ll get whole bunch of mesh (.bobj.gz) files after a simulation. One set for preview, and another for final. Each set has a .gz file for each frame of the animation. Each file contains the simulation result - so youâ← ll need them.

- Currently these files will not be automatically deleted, so it is a good idea to e.g. create a dedicated directory to keep simulation results. Doing a fluid simulation is similar to clicking the ANIM button - you currently have to take care of organizing the fluid surface meshes in some directory yourself. If you want to stop using the fluid simulation, you can simply delete all the `*fluid*.bobj.gz` files.

- Before running a high resolution simulation that might take hours, check the overall timing first by doing lower resolution runs.

- Fluid objects must be completely inside the bounding box of the domain object. If not, baking may not work correctly or at all. Fluid and obstacle objects can be meshes with complex geometries. Very thin objects might not appear in the simulation, if the chosen resolution is too coarse to resolve them (increasing it might solve this problem).

- Note that fluid simulation parameters, such as inflow velocity or the active flag can be animated with Fluidsim Ipos (see above).

- Donâ← t try to do a complicated scene all at once. Blender has a powerful compositor that you can use to combine multiple animations.

 For example, to produce an animation showing two separate fluid flows while keeping your domain small, render one .avi using the one flow. Then move the domain and render another .avi with the other flow using an alpha channel (in a separate B&W .avi?). Then, composite both .aviâ← s using the compositorâ← s add function. A third .avi is usually the smoke and mist and it is laid on top of everything as well. Add a rain sheet on top of the mist and spray and youâ← ll have quite a storm brewing! And then lightning flashes, trash blowing around, all as separate animations, compositing

the total for a truly spectacular result.

- If youâ← re having trouble, or something isnâ← t working as you think it should - let me know: send the .blend file and a problem description to `nils at thuerey dot de`. Please check these wiki pages and the blenderartists-forum before sending a mail!

Limitations & Workarounds

- One domain per blender file (as of Version 2.42), but you can have multiple fluid objects.

 Workaround: For previews, move the domain around to encompass each fluid flow part, and then for final, scale up the size of the domain to include all fluid objects (but computation will take longer). This is actually a benefit, because it lets you control how much compute time is used, by varying the size and location of the domain.

- If the setup seems to go wrong, make sure all the normals are correct (hence, enter Edit mode, select all, and recalculate normals once in a while).

- Currently thereâ← s a problem with zero gravity simulation - simply select a very small gravity until this is fixed.

- If an object is initialized as Volume, it has to be closed and have an inner side (a plane wonâ← t work). To use planes, switch to Shell, or extrude the plane.

- Blender freezes after clicking BAKE. Pressing Esc makes it work again after a while - this can happen if the resolution is too high and memory is swapped to hard disk, making everything horribly slow. Reducing the resolution should help in this case.

- Blender crashes after clicking BAKE - this can happen if the resolution is really high and more than 2GB are allocated, causing Blender to crash. Reduce the resolution. Many operating systems limit the total amount of memory that can be allocated by a *process*, such as Blender, even if the *machine* has more memory installed. Suxâ←¦

- The meshes should be closed, so if some parts of e.g. a fluid object are not initialized as fluid in the simulation, check that all parts of connected vertices are closed meshes. Unfortunately, the Suzanne (monkey) mesh in Blender is not a closed mesh (the eyes are separate).

- If the fluid simulation exits with an error message (stating e.g. that the â← init has failedâ←), make sure you have valid settings for the domain object, e.g. by resetting them to the defaults.

- To import a single fluid surface mesh you can use this script: .bobj.-Import-Script.

- You may not be able to bake a fluid that takes more than 1GB, not even with the LargeAddressAware build - it might be a limitation of the current fluid engine.

- Note that first frame may well take only a few hundred MBs of RAM memory, but latter ones go over one GB, which may be why your bake fails after awhile. If so, try to bake one frame at the middle or end at full res so youâ← ll see if it works.

- Memory used doubles when you set surface subdivision from 1 to 2.

- Using â← generate particlesâ← will also add memory requirements, as they increase surface area and complexity. Ordinary fluid-sim generated particles probably eat less memory.

TODO: check these links...

See Also

- Tutorial 1: Very Basic Introduction
- Tutorial 2: The Next Step
- Tutorial 1&2 Gui Changes for newer builds
- Another BSoD fluid tutorial
- Developer documentation (implementation, dependencies, â←¦)

External Documentation

- An Introduction to Fluid Simulations in Blender (video) (Blendernation link) - Learn the basics of how to set up a fluid simulation in Blender with an obstacle.
- Fluid Simulator Tutorial (video) (Blendernation link) - Very easy to understand video-tutorial to fluid simulation newcomers. Also covers some of the most common pitfalls.
- Guide on Blender Fluid Simulatorâ← s Parameters (Blendernation link)

Acknowledgements

The integration of the fluid simulator was done as a Google Summer-of-Code project. More information about the solver can be found at www.ntoken.com. These Animations were created with the solver before its integration into blender: Adaptive Grids, Interactive Animations. Thanks to Chris Want for organizing the Blender-SoC projects, and to Jonathan Merrit for mentoring this one! And of course thanks to Google for starting the whole thingâ←¦ SoC progress updates were posted here: SoC-Blenderfluid-Blog at PlanetSoC.

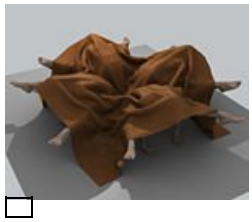
The solver itself was developed by Nils Thuerey with the help and supervision of the following people: U. Ruede, T. Pohl, C. Koerner, M. Thies, M. Oechsner and T. Hofmann at the Department of Computer Science 10 (System Simulation, LSS) in Erlangen, Germany.

ttp://www10.informatik.uni-erlangen.de/~sinithue/img/lsslogo.png http://www10.informatik.uni-erlangen.de/~sinithue/img/unierlanger

Cloth Simulation



Cloth example.



Cloth on carved wooden men (made by motorsep).



Cloth example.

Cloth simulation is one of the hardest aspects of CG, because it is a deceptively simple real-world item that is taken for granted, yet actually has very complex internal and environmental interactions. After years of development, Blender has a very robust cloth simulator that is used to make clothing, flags, banners, and so on. Cloth interacts with and is affected by other moving objects, the wind and other forces, as well as a general aerodynamic model, all of which is under your control.

Description

Mode: Object Mode

Panel: Editing context â Modifiers panel

Hotkey: F7 to get to Object context; repeat to change sub-context.

A piece of cloth is any mesh, open or enclosed, that has been designated as cloth. The Cloth panels are located in the Physics sub-context and consist of three panels of options. Cloth is either an open or closed mesh and is mass-less, in that all cloth is assumed to have the same density, or mass per square unit.

Cloth is commonly modeled as a mesh grid primitive, or a cube, but can also be, for example, a teddy bear. However, Blenderâ ← s Softbody system provides better simulation of closed meshes; Cloth is a specialized simulation of fabrics.

Once the object is designated as Cloth, a Cloth modifier will be added to the objectâ ← s modifier stack automatically. As a modifier then, it can interact with other modifiers, such as Armature and Smooth. In these cases, the ultimate shape of the mesh is computed in accordance with the order of the modifier stack. For example, you should smooth the cloth *after* the modifier computes the shape of the cloth.

So you edit the Cloth in two places. In the F7 Physics buttons to edit the properties of the cloth and in the Modifier stack where you edit the Modifier properties related to display and interaction with other modifiers.

You can Apply the cloth modifier to freeze, or lock in, the shape of the mesh at that frame, which removes the modifier. For example, you can drape a flat cloth over a table, let the simulation run, and then apply the modifier. In this sense, you are using the simulator to save yourself a lot of modeling time.

Results of the simulation are saved in a cache, so that the shape of the mesh, once calculated for a frame in an animation, does not have to be recomputed again. If changes to the simulation are made, the user has full control over clearing the cache and re-running the simulation. Running the simulation for the first time is fully automatic and no baking or separate step interrupts the workflow.

Computation of the shape of the cloth at every frame is automatic and done in the background; thus you can continue working while the simulation is computed. However it is CPU-intensive and depending on the power of your PC and the complexity of the simulation, the amount of CPU needed to compute the mesh varies as does the lag you might notice.

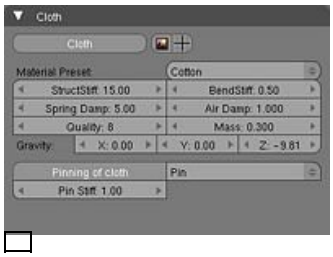
Donâ← t jump ahead
If you set up a cloth simulation but Blender has not computed the shapes for the duration of the simulation, and if you jump ahead a lot of frames forward in your animation, the cloth simulator may not be able to compute or show you an accurate mesh shape for that frame, if it has not previously computed the shape for the previous frame(s).

Workflow

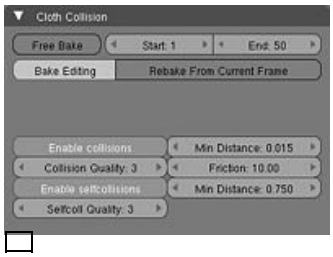
A general process for working with cloth is to:

1. Model the cloth object as a general starting shape.
2. Designate the object as a â← clothâ← in the Physics sub-context of the Object context (F7).
3. Model other deflection objects that will interact with the cloth. Ensure the Deflection modifier is last on the modifier stack, after any other mesh deforming modifiers.
4. Light the cloth and assign materials and textures, UV-unwrapping if desired.
5. If desired, give the object particles, such as steam coming off the surface.
6. Run the simulation and adjust Options to obtain satisfactory results. The timeline windowâ← s VCR controls are great for this step.
7. Optionally age the the mesh to some point in the simulation to obtain a new default starting shape.
8. Make minor edits to the mesh on a frame-by-frame basis to correct minor tears.

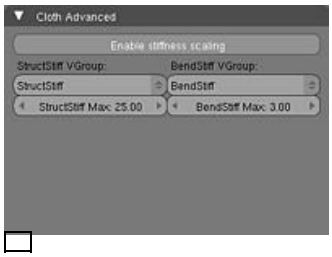
Cloth Panel Options



Cloth panel.



Cloth Collisions panel.



Cloth Advanced functions panel.

Cloth is controlled by the three panels shown above. Each of the options (buttons and controls) on those three panels is fully documented in the Reference manual. This section discusses how to use those options to get the effect you want.

Cloth Definition

First, enable Cloth. Set up for the kind of cloth you are simulating. You can choose one of the presets to have a starting point.

As you can see, the heavier the fabric, the more stiff it is and the less it stretches and is affected by the air.

Next, position and pin the cloth in the desired position. Pinning is described below. Note that if you move the cloth object *after* you have already run some simulations, you must unprotect and clear the cache; otherwise, Blender will use the position of the current/cached meshâ← s vertices when trying to represent where they are. Editing the shape of the mesh, after simulation, is also discussed below. You may disable the cloth and edit the mesh as a normal mesh editing process.

Finally, use the Timeline window Play button, or press AltA in the 3D View to run the simulation. Your cloth will fall and interact with Deflection objects as it would in the real world.

Using Simulation to Shape/Sculpt a Mesh

You can Apply the Cloth modifier at any point to freeze the mesh in position at that frame. You can then re-enable cloth, setting the start and end frames from which to run the simulation forward.

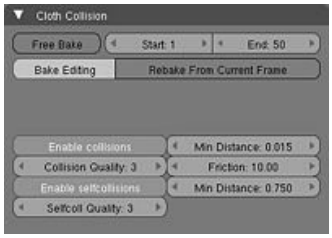
Another example of aging is a flag. Define the flag as a simple grid shape and pin the edge against the flagpole. Simulate for 50 frames or so, and the flag will drop to its â← restâ← position. Apply the Cloth modifier. If you want the flag to flap or otherwise move in the scene, re-enable it for the frame range when it is in camera view.

Collisions

In most cases, a piece of cloth does not just hang there in 3D space, it collides with other objects in the environment. To ensure proper simulation,there are several items that have to be set up and working together:

1. The Cloth object must be told to participate in Collisions.
2. Optionally (but recommended) tell the cloth to collide with itself.
3. Other objects must be visible to the Cloth object *via* shared layers.
4. The other objects must be mesh objects.
5. The other objects may move or be themselves deformed by other objects (like an armature or shape key).
6. The other mesh objects must be told to deflect the cloth object.
7. The blend file must be saved in a directory so that simulation results can be saved.
8. You then Bake the simulation. The simulator computes the shape of the cloth for a frame range.
9. You can then edit the simulation results, or make adjustments to the cloth mesh, at specific frames.
10. You can make adjustments to the environment or deforming objects, and then re-run the cloth simulation from the current frame forward.

Cloth Collision



Cloth Collisions panel.

Now you must tell the Cloth object that you want it to participate in collisions. For the cloth object, locate the Cloth Collision panel, shown to the right:

Enable Collisions
: LMB click this to tell the cloth object that it needs to move out of the way.

Min Distance

As another object gets this close to it (in Blender Units), the simulation will start to push the cloth out of the way.

Collision Quality
: A general setting for how fine and good a simulation you wish. Higher numbers take more time but ensure less tears and penetrations through the cloth.

Friction
: A coefficient for how slippery the cloth is when it collides with the mesh object. For example, silk has a lower coefficient of friction than cotton.

Self-collisions

Real cloth cannot permeate itself, so you normally want the cloth to self-collide.

Enable Selfcollisions
: LMB click this to tell the cloth object that it should not penetrate itself. This adds to simulation compute time, but provides more realistic results. A flag, viewed from a distance does not need this enabled, but a close-up of a cape or blouse on a character should have this enabled.

Min Distance
: If you encounter problems, you could also change the Min Distance value for the self-collisions. The best value is 0.75, for fast things you better take 1.0. The value 0.5 is quite risky (most likely many penetrations) but also gives some speedup.

Selfcoll Quality
: For higher self-collision quality just increase the Selfcoll Quality and more selfcollision layers can be solved. Just keep in mind that you need to have at least the same Collision Quality value as the Selfcoll Quality value.

Regression blend file: Cloth selfcollisions.

Shared Layers

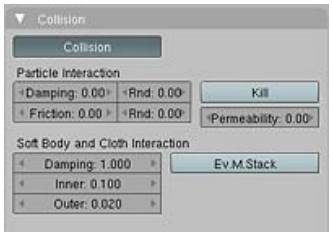
For example, suppose you have two objects: a pair of Pants on layers 2 and 3, and your Character mesh on layers 1 and 2. You have enabled the Pants as cloth as described above. You must now make the Character â← visibleâ← to the Cloth object, so that as your character bends its leg, it will push the cloth. This principle is the same for all simulations; simulations only interact with objects on a shared layer. In this example, both objects share layer 2.

To view/change an objectâ← s layers, RMB click to select the object in Object mode in the 3D view. M to bring up the â← Move Layersâ← popup, which shows you all the layers that the object is on. To put the object on a single layer, LMB click the layer button. To put the object on multiple layers, â § ShiftLMB the layer buttons. To remove an object from a selected layer, simply â § ShiftLMB the layer button again to toggle it.

Mesh Objects Collide

If your colliding object is not a mesh object, such as a NURBS surface, or text object, you must convert it to a mesh object. To do so, select the object in object mode, and in the 3D View header, select Object â Convert Object Type (AltC), and select Mesh from the popup menu.

Cloth - Object collisions

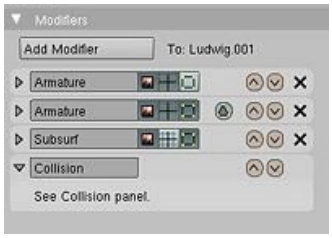


Collision settings.
The cloth object needs to be deflected by some other object. To deflect a cloth, the object must be enabled as an object that collides with the cloth object. To enable Cloth - Object collisions, you have to enable deflections on the collision object (not on the cloth object).

In the Buttons window, Object context, Physics sub-context, locate the Collision panel shown to the right. It is also important to note that this collision panel is used to tell all simulations that this object is to participate in colliding/deflecting other objects on a shared layer (particles, soft bodies, and cloth).

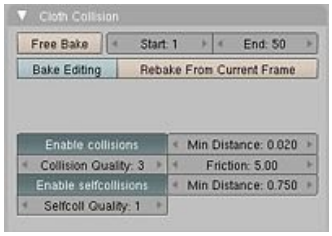
Beware
There are three different Collision panels, all found in the Physics sub-context. The first (by default), a tab beside the Fields panel, is the one needed here. The second panel, a tab in the Soft Body group, concern softbodies (and so has nothing to see with clothes). And we have already seen the last one, by default a tab beside the Cloth panel.

Mesh Object Modifier Stack



Collision stack.
The objectâ← s shape deforms the cloth, so the cloth simulation must know the â← trueâ← shape of that mesh object at that frame. This true shape is the basis shape as modified by shape keys or armatures. Therefore, the Collision modifier must be **after** any of those. The image to the right shows the Modifiers panel for the Character mesh object (not the cloth object).

Bake Collision



After you have set up the deflection mesh for the frame range you intend to run the simulation (including animating that mesh *via* armatures), you can now tell the cloth simulation to compute (and avoid) collisions. Select the cloth object and in the Object context, Physics sub-context, set the Start and End settings for the simulation frames you wish to compute, and click the Bake button.

Start
: The starting frame of the animation when you want the cloth to start responding.

End
: The end frame of the simulation. The cloth will remain â← frozenâ← after End.

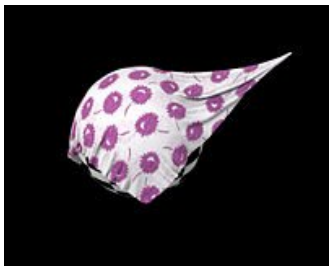
Bake
: Starts the simulation process.

You cannot change Start or End without clearing the bake simulation. When the simulation has finished, you will notice you have the option to free the bake, edit the bake and re-bake:

Free Bake
: Deletes the simulation, enabling you to make changes and then start over.

Bake Editing
: Enables you to protect the cache and prevent the simulation from recomputing the mesh shape. You can then edit the mesh shape manually.

Rebake From Current Frame
: Enables you to re-compute the cloth simulation from the current frame to End.



□
Standard settings cloth collide with sphere.
Thereâ← s a few things youâ← ll probably notice right away. First, it will bake significantly slower than before, and it will probably clip through the box pretty bad as in the picture on the right.

Editing the cached simulation

The cache contains the shape of the mesh at each frame. You can edit the cached simulation, when you baked the simulation and pressed the Bake Editing button. Just go to the frame you want to fix and â Tab into Edit mode. There you can move your vertices using all of Blenderâ← s mesh shaping tools. When you exit, the

shape of the mesh will be recorded for that frame of the animation. If you want Blender to resume the simulation using the new shape forward, LMB click â← Rebake from next Frame and play the animation. Blender will then pick up with that shape and resume the simulation.

Edit the mesh to correct minor tears and places where the colliding object has punctured the cloth.

If you add, delete, extrude, or remove the vertices in the mesh, Blender will take the new mesh as the starting shape of the mesh back to the *first frame* of the animation, replacing the original shape you started with, up to the frame you were on when you edited the mesh. Therefore, if you change the content of a mesh, when you â Tab out of Edit mode, you should unprotect and clear the cache so that Blender will make a consistent simulation.

Troubleshooting

If you encounter some problems with collision detection there are two ways to fix them:

- The fastest solution would be to put up the Min Distance setting under the Cloth Collision panel. This will be the fastest way to fix the clipping, however, it will be less accurate and wonâ← t look as good. Using this method tends to make it look like the cloth is resting on air, and gives it a very rounded look.

- A second method is to increase the Quality (in the first Cloth panel). This results in smaller steps for the simulator and therefore to a higher probability that fast moving collisions get catch. You can also increase the Collision Quality to perform more iterations to get collisions solved.

- If none of the methods help, you can easily edit the cached/baked result in Edit mode afterwards.

- My Cloth is torn by the deforming mesh - he â← Hulks Outâ← : Increase its structural stiffness (StructStiff setting, Cloth panel) very high, like 1000.

Subsurf modifier
A bake/cache is done for every subsurf level so please use **one equal** subsurf level for render and preview.

Examples

To start with cloth, the first thing you need, of course, is some fabric. So, lets delete the default cube and add a plane. I scaled mine up along the Y axis, but you donâ← t have to do this. In order to get some good floppy and flexible fabric, youâ← ll need to subdivide it several times. I did it 8 times for this example. So â Tab into Edit mode, and press W â Subdivide multi, and set it to 8.

Now, weâ← ll make this cloth by going to the Object context (F7) â Physics sub-context. Scroll down until you see the Cloth panel, and press the Cloth button. Now, a lot of settings will appear, most of which weâ← ll ignore now.

Thatâ← s all you need to do to set your cloth up for animating, but if you hit AltA, your lovely fabric will just drop very un-spectacularly. Thatâ← s what weâ← ll cover in the next two sections about pinning and colliding.

Pinning of cloth



Cloth in action.
The first thing you need when pinning cloth are Vertex Groups. There are several ways of doing this including using the Weight Paint tool to paint the areas you want to pin (see the Weight paint section of the manual).

Once you have a vertex group set, things are pretty straightforward, all you have to do is press the Pinning of cloth button in the Cloth panel and select which vertex group you want to use, and the stiffness you want it at. You can leave the stiffness as it is, default value is fine.

Smoothing of Cloth

Now, if you followed this from the previous section, your cloth is probably looking a little blocky. In order to make it look nice and smooth like the picture you need to apply a Smooth and/or Subsurf modifier in the Modifiers panel under the Editing context (F9). Then, in the same context, find the Links and Materials panel (the same one you used for vertex groups) and press Set Smooth.

Now, if you hit AltA, things are starting to look pretty nice, donâ← t you think?

Cloth on armature

Cloth deformed by armature and also respecting an additional collision object: Regression blend file.

Using Cloth for Softbodies



Using cloth for softbodies.
Cloth can also be used to simulate softbodies. Itâ← s for sure not itâ← s main purpose but it works nonetheless. The example image uses standard Rubber material, no fancy settings, just AltA.

Blend file for the example image: Using Cloth for softbodies.

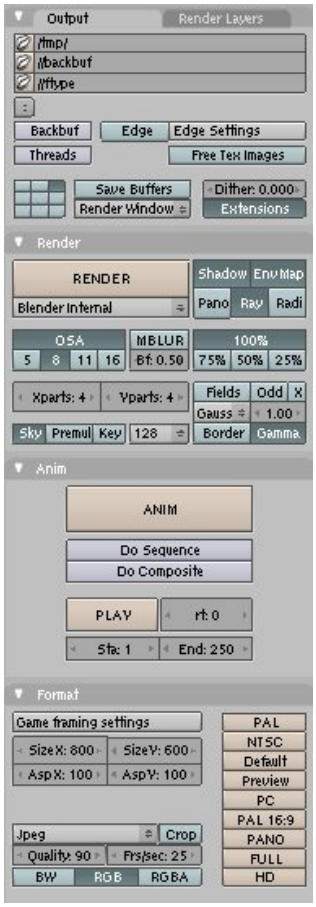
Cloth with Wind



Flag with wind applied.
Regression blend file for Cloth with wind and selfcollisions (also the blend for the image above): Cloth flag with wind and selfcollisions.

Rendering is the final process of CG (short of post processing, of course) and is the phase in which the image corresponding to your 3D scene is finally created. Rendering is a CPU-intensive process. You can render an image on your computer, or use a render farm which is a network of PC's that each work on a different section of the image or on different frames. This section provides a full explanation of the features in Blender related to the process of producing your image or animation.

Introduction



Rendering Buttons.

The rendering buttons window is accessed via the Scene Context and Render Sub-context (F10 or the  button). The rendering Panels and Buttons are shown in *Rendering Buttons.*.

These buttons are organized into panels, which are:

- Output - controls the output of the render pipeline
- Render Layers - controls which layers and passes to render
- Render - controls the actual rendering process of a still shot
- Anim - controls the rendering of a series of frames to produce an animation
- Bake - pre-computes certain aspects of a render
- Format - controls the format and encoding of the picture or animation
- Stamp - stamps the frames with identifying and configuration control item information

Tabs

To save screen space, some of the panels may be tabbed under another; for example, the Layers panel is a tab folder under the Output panel. To reveal it, simply click the tab header.

Yafray

If you have installed Yafray, options to control it will appear as two tabs under the Render panel once you have selected it as a rendering engine

Overview

The rendering of the current scene is performed by pressing the big RENDER button in the Render panel, or by pressing F12 (you can define how the rendered image is displayed on-screen in the Render Output Options). See also the Render Window.

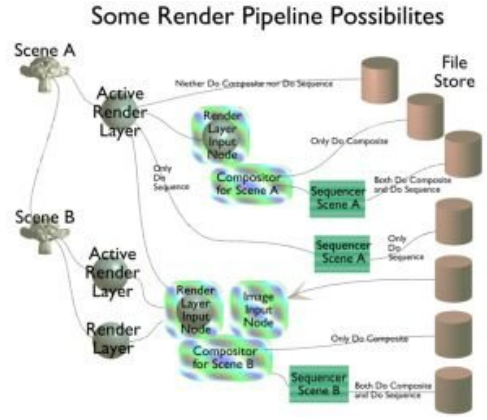
A movie is produced by pressing the big ANIM animation button in the Anim panel. The result of a rendering is kept in a buffer and shown in its own window. It can be saved by pressing F3 or via the File->Save Image menu using the output option in the Output panel. Animations are saved according to the format specified, usually as a series of frames in the output directory. The image is rendered according to the dimensions defined in the Format Panel (*Image types and dimensions.*).

Workflow

In general, the process for rendering is:

1. Create all the objects in the scene
2. Light the scene and position the camera
3. Render a test image at 25% or so without oversampling or raytracing etc. so that it is very fast and does not slow you down
4. Set and Adjust the materials/textures and lighting
5. Iterate the above steps until satisfied at some level of quality
6. Render progressively higher-quality full-size images, making small refinements and using more compute time
7. Save your images

Render Workbench Integration



Blender has three independent rendering workbenches which flow the image processing in a pipeline from one to the other in order:

- Rendering Engine
- Compositor
- Sequencer

You can use each one of these independently, or in a linked workflow. For example, you can use the Sequencer by itself to do post processing on a video stream. You can use the Compositor by itself to perform some color adjustment on an image. You can render the scene, via the active Render Layer, and save that image directly, with the scene image computed in accordance with the active render layer, without using the

Compositor or Sequencer. These possibilities are shown in the top part of the image to the right.

You can also link scenes and renders in Blender as shown, either directly or through intermediate file storage. Each scene can have multiple render layers, and each Render Layer is mixed inside the Compositor. The active render layer is the render layer that is displayed and checked active. If the displayed render layer is not checked active/enabled, then the next checked render layer in the list is used to compute the image. The image is displayed as the final render if Do Composite and Do Sequence are NOT enabled.

If Do Composite is enabled, the render layers are fed into the Compositor. The noodles manipulate the image and send it to the Composite output, where it can be saved, or, if *Do Sequence* is on, it is sent to the Sequencer.

If Do Sequence is enabled, the result from the compositor (if Do Composite is enabled) or the active Render layer (if Do Composite is not enabled) is fed into the Scene strip in the Sequencer. There is is manipulated according to the VSE settings, and finally delivered as the image for that scene.

Things get a little more complicated when a .blend file has multiple scenes, for example Scene A and Scene B. In Scene B, if Do Composite is enabled, the Render Layer node in Scene B's compositor can pull in a Render Layer from Scene A. Note that this image will not be the post-processed one. If you want to pull in the composited and/or sequenced result from Scene A, you will have to render Scene A out to a file using Scene A's compositor and/or sequencer, and then use the Image input node in Scene B's compositor to pull it in.

The bottom part of the possibilities graphic shows the ultimate blender: post-processed images and a dynamic component render layer from Scene A are mixed with two render layers from Scene B in the compositor, then sequenced and finally saved for your viewing enjoyment.

These examples are only a small part of the possibilities in using Blender. Please read on to learn about all the options, and then exercise your creativity in developing your own unique workflow.

Sequencer from/to Compositor

To go from the Compositor to the Sequencer, enable both "Do Sequence" and "Do Composite". In the Compositor, the image that is threaded to the Composite Output node is the image that will be processed in the Scene strip in the VSE.

The way to go from the Sequencer to the Compositor is through a file store. Have one scene "Do Sequence" and output to an image sequence or mov/avi file. Then, in another scene, "Do Composite" and using an image input node to read in the image sequence or mov/avi file.

Render Window Options

Once you press F12 or click the big Render button, your image is computed and display begins. Depending on the Output Panel Option, the image is shown in a separate Render Window, Full Screen, or in a UV/Image Editor window.

You can render a single frame, or many frames in an animation. You can cancel the rendering by pressing Esc. If rendering a sequence of frames, the lowest number frame is rendered, then the next, and so on in increasing frame number sequence.

The Render Window can be invoked in several ways:

Render Panel->Render or F12

renders the current frame, as seen by the active camera, using the selected renderer (Blender Internal or Yafray)

3D View Window Header->LMB OpenGL Render button (far right)
Renders a quick preview of the 3D View window

Anim Panel->Anim CtrlF12
Render the Animation from the frame start to and included in the End frame, as set in the Anim panel.

3D View Window Header->CtrlLMB OpenGL Render button (far right)
Renders a quick animation of the 3D View window using the fast OpenGL render

Output Options needs to be set correctly. In the case of Animations, each frame is written out in the Format specified.

Rendering the 3D View Animation using the OpenGL is useful for showing armature animations.

Showing Previous Renders

If the Blender Internal render was used to compute the image, you can look at the previous render:

Render->Show Render Buffer
F11 - Pops up the Render Window and shows the last rendered image (even if it was in a previously opened & rendered .blend file).

Render->Play Back Rendered Animation
CtrlF11 - Similar as for the single frame, but instead plays back all frames of the rendered animation.

Render Window usage

Once rendering is complete and the render window is active, you can:

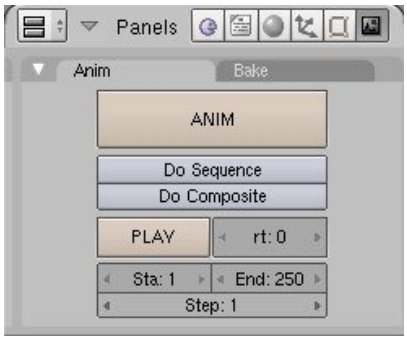
- A - Show/hide the alpha layer.
- Z - Zoom in and out. Pan with the mouse. You can also mousewheel to zoom
- J - Jump to other Render buffer. This allows you to keep more than one render in the render window, which is very useful for comparing subtleties in your renders. How to use it:

1. Press Render or F12
2. Press J to show the empty buffer (the one we want to "fill" with the new image)
3. Go back to the Blender modeling window. You can send the render window to the background by pressing Esc. *Do not* close, or minimize the render window!
4. Make your changes.
5. Render again.
6. Press J to switch between the two renderings.

- LMB - Clicking the left mouse button in the Render window displays information on the pixel underneath the mouse pointer. The information is shown in a text area at the bottom left of the Render output window. Moving the mouse while holding the LMB will dynamically update the information. This information includes:
 - Red, Green, Blue and Alpha values
 - Distance in Blender Units from the camera to the pixel under the cursor (Render Window only)

For Alpha values to be computed, RGBA has to be enabled. Z-depth (distance from camera) if computed only if Z is enabled on the Render Layers tab.

Step Render Frame



Step Frame Option.
Blender allows you to do faster animation renders skipping some Frames. You can set the step in the Render Panel as you can see in the picture.

Once you have your video file rendered, you can play it back in the real speed (fps). For that you just need to set the same Step parameter you used to render it.

If you play the stepped video in a normal speed in an external player, the general speed will be *Step* times faster than a video rendered with all the frames (eg. a video with 100 frames rendered with Step 4 @ 25 fps will have only 1 second. The same video rendered with Step 1 (the default value) will have 4 seconds of length).

If you want to use this parameter for rendering through command line you can use -j STEP, where STEP stands for the number of steps you want to use.

If you want to use this parameter for playing video files through the command line, you need to use the parameter -j STEP following the -a (which stands for the playback mode).

./blender -a -s 1 -e 100 -p 0 0 -f 25 1 -j 4 "//videos/0001.jpg"

- Rendered stepped frames with output video format such as FFMpeg (always) or AVI Jpeg (sometimes) produce a corrupted video (missing end frames). in Blender 2.48a. Therefore the rendered video can't be played-back properly.

Therefore I suggest to work with the video output format of JPG, it works fine all the time.

Distributed Render Farm

There are several levels of CPU allocation that you can use to decrease overall render time by applying more brainpower to the task.

First, if you have multi-core CPU, you can increase the number of threads, and Blender will use that number of CPUs to compute the render.

Second, if you have a local area network with available PC's, you can split the work up by frames. For example, if you want to render a 200 frame animation, and have 5 PC's of roughly equal processing power, you can allocate PC#1 to produce frames 1-40, PC#2 to frames 41-80, and so on. If one PC is slower than the others, simply allocate fewer frames to that PC. To do LAN renders, map the folder containing the .blend file (in which you should have packed your external data, like the textures, â←¦) as a shareable drive. Start Blender on each PC and open the .blend file. Change the Start and End frame counts on that PC, but do not save the .blend file. Start Rendering. If you use relative paths for your output pathspec, the rendered frames will be placed on the host PC.

Third, you can do WAN rendering, which is where you email or fileshare or Verse-share the .blend file (with packed datas!) across the Internet, and use anyone's PC to perform some of the rendering. They would in turn email you the finished frames as they are done. If you have reliable friends, this is a way for you to work together.

Fourth, you can use a render farm service. This service, like BURP, is run by an organization. You email them your file, and then they distribute it out across their PC's for rendering. BURP is mentioned because it is free, and is a service that uses fellow BlenderHead PC's with a BOINC-type of background processing. Other services are paid subscription or pay-as-you-go services.

Show the Render Window

The Render window can be invoked/shown in 2 ways:

Rendering

Output Options needs to be set correctly. See Rendering for details on rendering itself.

Render â Render Current Frame (F12)
: As the title already says, it renders the current frame and shows the Render window.

Render â Render Animation (CtrlF12)
: Same for Animation.

Showing

Render â Show Render Buffer (F11)
: Pops up the Render window and shows the last rendered image (even if it was in a previously opened & rendered .blend file).

Render â Play Back Rendered Animation (CtrlF11)
: Similar as for the single frame, but instead plays back all frames of the rendered animation.

Render Window usage

- A - Show/hide the alpha layer.
- Z - Zoom in and out. Pan with the mouse.
- J - Render buffer switch. This allows you to keep two images in the render window, which is very useful for comparing subtleties in your renders. How to use it:

 1. Press RENDER or F12.
 2. Press J to show the empty buffer (the one we want to â← fillâ← with the new image). Note that you can keep one image in a buffer and compare it with subsequent renders.
 3. Make your changes.
 4. Render again.
 5. Press J to switch between the two renderings.

- Esc - Cancel a running render process
- LMB - Clicking the left mouse button in the Render window displays information on the pixel underneath the mouse pointer. The information is shown in a text area at the bottom left of the Render output window. Moving the mouse while holding the LMB will dynamically update the information. This information includes:
 - ♦ Red, Green, Blue and Alpha values.
 - ♦ Distance in Blender Units from the camera to the pixel under the cursor.

See also

- Rendering
- Render Output Options

We know that around the world, our users have PC's of widely varying power. Rendering is *the* process in CG that can chew up CPU and disk space like no tomorrow. Especially in corporate environments, it is easy to fill up terabyte servers by uploading ten hour-long DV tapes and doing some editing. So, there are lots of options try to shoehorn a big job into a small PC by providing you with multiple sets of options that chunk up the work as best we can, while still preserving image integrity.

This page discusses the main options found on the Render panel, and subsequent pages give you more.

Shadow

Enable this button to compute and render shadows. Shadows are cast by shadow-casting lights and are received by shadow-able materials. The results of this calculation are made available in the Shadow render pass. For more information on lights, materials, or render passes, please consult the application sections of the wiki.

EnvMap

Environmental mapping changes the background and color of objects based on the light cast by the world environment maps. Enable for more realistic rendering.

Raytracing

Raytracing is a more precise method of computing the color of a surface, especially reflective surfaces. It is enabled by clicking the Ray button.

Octree resolution

Mode: All Modes

Panel: Render Context → Render

Hotkey: F10

When raytracing a really big scene with small objects, render times can grow exponentially. To help keep render times down (but use more memory), you can increase the Octree memory allocation. When Raytracing is enabled (Ray button, next to Render button on the Render panel), the Octree select appears at the bottom of the panel, with selectable settings ranging up to 512.

The OctTree is a ray-tracing acceleration structure. It subdivides the whole scene into a regular 3D grid of cells and places every polys of every objects into their corresponding cell. When rendering, it is quicker to determine which cell a ray is intersecting and then test the polys in those cells instead of testing every polys in the scene.

When you have a large low-poly scene with a small high-poly model in the middle, if the octtree resolution is too small, most of the scene's poly that comes from the small high-poly model, end up in a very small number of cells. And so, even though it is relatively quick to figure which cell to test, when a cell contains a large number of polys, we gained nothing because all those polys must be tested.

On the other hand, if the octtree resolution is too large for a given scene, then the raytracer looses time checking for cells that contains no polys. That is why each scene have its own octtree resolution sweet spot that must be found experimentally.

So, use a high setting for a large open space with small areas that have high-poly structures, and low settings for scenes where the polys (faces) are evenly distributed. For more information, consult the Release Notes for Octree resolution

Radiosity

When light hits an object, some color spectrum of the light is absorbed and other colors are reflected to our eyes. Radiosity is when that light also hits an object close to it, giving color to it, which then in turn is sent to our eyes. Sometimes called color bleeding, because the color of one object bleeds onto the one next to it, radiosity gives much more photorealism.

Percentage Size

Calculating a full size image takes time. For interim renders, Click the 75%, 50% or 25% boxes to render a smaller image (which takes less time). When looking at the render window, you can always expand it and roll your mousewheel to zoom in / expand the image.

Parts Rendering

Mode: All Modes

Panel: Render Context → Render

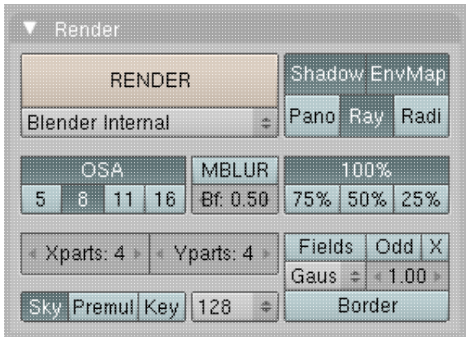
Hotkey: F10

Description

It is possible to render an image in pieces, one after the other, rather than all at one time. This can be useful for very complex scenes, where rendering small sections one after the other only requires computation of a small part of the scene, which uses less memory.

On a multi-core CPU, each part is assigned to a CPU/core, so setting up multiple parts, coupled with increasing the number of threads, speeds up render time by using all of your PC's processing power.

Options



Rendering by parts buttons (F10).
By setting values different from 1 in the Xparts and Yparts NumButtons in the Render Panel (*Rendering by parts buttons.*), you force Blender to divide your image into a grid of Xparts times Yparts sub-images, which are then rendered one after the other and finally assembled together.

Memory is allocated per thread; with each thread doing a full tile render. Almost all geometry (render faces/vertices) is checked for each tile it renders, giving some extra overhead. There's also thread tables for AO and area-shadow lamp sampling, precalculated and allocated in advance. Lastly; there's still quite some locking going on, for each memory allocation call to the operating system (malloc) for example.

A quick rule-of-thumb is to make sure the total number of threads renders less than one quarter of the entire image, to prevent too much memory allocated:

8 threads: use 64 tiles (for example, X and Y Parts = 8) 16 threads: use 128 tiles 128 threads: use 1024 tiles

Whether such giant amount of tiles (1024) gives unacceptable overhead is unknown, but quite likely. If you use heavy raytracing on relatively simple scenes it might work though. Again, if the tiles are square in terms of pixels, then Save Buffers can be used to ease memory.

Limitations & Work-arounds

Blender cannot handle more than 64 parts in the Y direction

Panoramic Rendering

Mode: All Modes

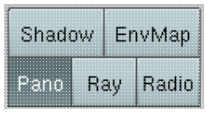
Panel: Render Context → Render

Hotkey: F10

Description

To obtain nice panoramic renderings, up to 5 times (!) a full 360Â° view of the horizon, Blender provides an automatic procedure.

Options



"Pano" Button.
You can, by decreasing the focal length of your camera, get a wider field of view, up to 173Â° (length of 1mm), but at cost of huge distortions in the image ("fish-eye" effect); furthermore, you won't be able to get wider than these 173Â°.

But Blender is able to render an image showing a 1800Â° panorama (5 full rotations) of the scene, as if the camera was rotating around its Y axis, with few distortions. For rendering a "real" panorama, enable the Pano button. Henceforth, the behaviour of some render and camera settings are changed:

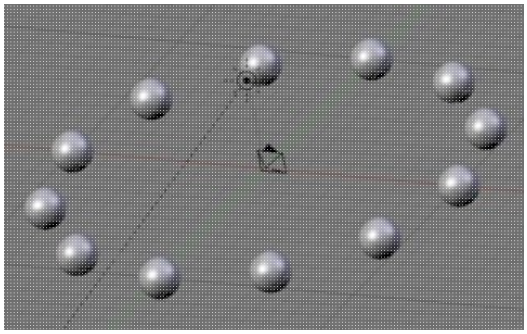
Camera
Lens:
: A 5 (mm) lens setting gives a 360Â° pano. The horizontal field of view is now proportional to this setting: 10mm gives a 180Â° pano, 2.5mm gives a 720Â° pano (two turns), 1mm gives a 1800Â° pano (5 turns), etc...
This change only affects the *horizontal* field of view: the *vertical* one behaves as usual (i.e. it is locked to 173Â° at maximum!). This means that if you want to render a *vertical* pano, you have to lay the camera on its side.

Rendering
Xparts
: Defines the number of "shots" aligned side by side: at minimum to 10 if you want a "correct" pano; the higher it is, the better is the result (lower distortions); the max number of shots is the width of the rendered picture, in pixels, divided by eight.
Yparts
: Its behaviour **isn't changed** from a "standard" rendering.
SizeX, SizeY
: As long as SizeX > SizeY, the *horizontal* field of view stays the same, as defined by Lens: (e.g., for a 5mm lens, 360Â°): the vertical field of view is proportional to the ratio height/width.
If SizeX < SizeY, the *vertical* field of view is locked to its maximum (173Â° for a Lens: of 1mm, 145Â° for a Lens: of 5mm, etc.): the horizontal field of view is proportional to the ratio width/height

(e.g., for a rendered picture twice as high as wide, and a Lens: of 5, we have an horizontal 180Â° pano, rather than a 360Â° one).

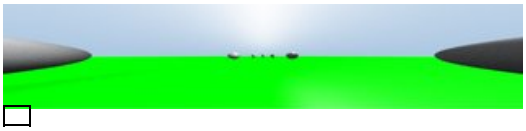
Examples

All this is quite complex, so here are some examples, all based on the same scene, to try to clarify it:



Test scene.

Examples of non-panoramic rendering



Non-panoramic rendering with Lens: = 1mm (horizontal field of view: 173Â°).



Non-panoramic rendering with Lens: = 5mm (horizontal field of view: 145Â°).



Non-panoramic rendering with Lens: = 10mm (horizontal field of view: 116Â°).

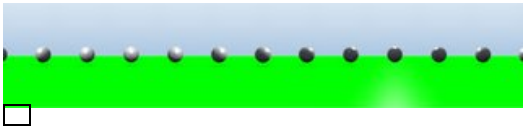
Examples of panoramic rendering



Panoramic rendering with Lens: = 5mm, and Xparts = 1. There's no difference from a rendering with the same lens without the Pano option!



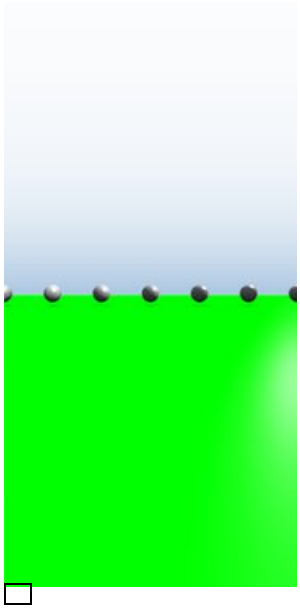
Panoramic rendering with Lens: = 5mm, and Xparts = 5. The distortions are still obvious, and the horizontal field of view is not yet full 360Â°...





Panoramic rendering with Lens: = 5mm, and Xparts = 10. Nearly perfect.

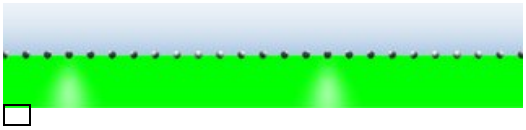
Panoramic rendering with Lens: = 5mm, and Xparts = 90. Compare with the previous one: very few differences...



Rendu panoramique avec Lens: = 5mm, and Xparts = 90, and twice as high as wide: we have a 180Â° pano instead of a 360Â° one, with a vertical field of view of 145 Â°!



Panoramic rendering with Lens: = 1mm, and Xparts = 90. Five complete turns (very useful!).



Panoramic rendering with Lens: = 2.5mm, and Xparts = 90. Two complete turns (very useful!).



Panoramic rendering with Lens: = 10mm, and Xparts = 90. 180Â° pano.

Author's note
Everything above about the panoramic rendering is written from my Blender user's experience: I've never looked at its renderer code...

Motion Blur Rendering

Mode: All Modes

Panel: Render Context → Render

Hotkey: F10

Description

Blender's animations are by default rendered as a sequence of *perfectly still* images. This is unrealistic, since fast moving objects do appear to be 'moving', that is, blurred by their own motion, both in a movie frame and in a photograph from a 'real world camera'. To obtain such a Motion Blur effect, Blender can be made to render the current frame and some more frames, in between the real frames, and merge them all together to obtain an image where fast moving details are 'blurred'.

Options



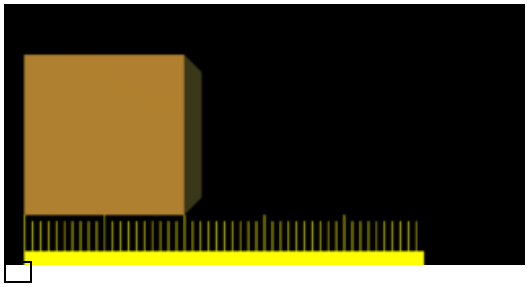
Motion Blur buttons (F10).

MBLUR
: Enables multi-sample motion blur. This makes Blender render as many 'intermediate' frames as the oversampling number is set to (5, 8, 11 or 16) and accumulate them, one over the other, on a single frame. The number-button Bf: or Blur Factor defines the length of the shutter time as will be shown in the example below. Setting the OSA Button is unnecessary since the Motion Blur process adds some antialiasing anyway, but to have a really smooth image OSA can be activated too. This makes each accumulated image have anti-aliasing (becareful with the render time!).

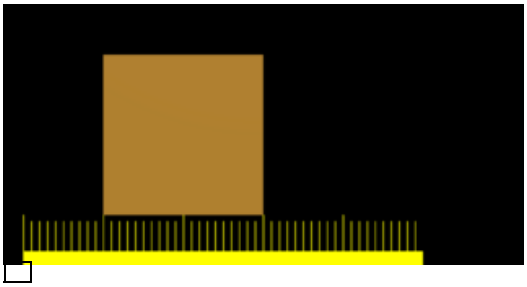
Examples

To better grasp the concept let's assume that we have a cube, uniformly moving 1 Blender unit to the right at each frame. This is indeed fast, especially since the cube itself has a side of only 2 Blender units.

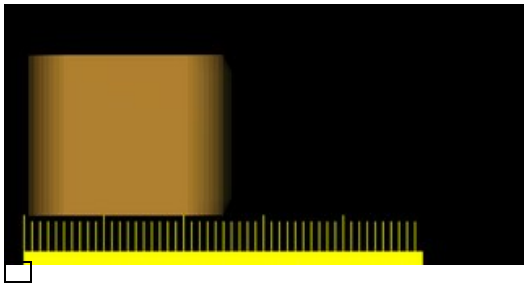
Frame 1 of moving cube without Motion Blur. shows a render of frame 1 without Motion Blur, *Frame 2 of moving cube without Motion Blur.* shows a render of frame 2. The scale beneath the cube helps in appreciating the movement of 1 Blender unit.



Frame 1 of moving cube without Motion Blur.



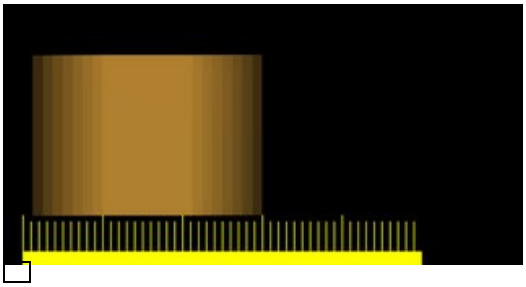
Frame 2 of moving cube without Motion Blur.



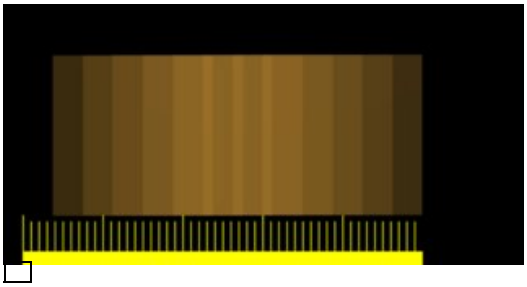
Frame 1 of moving cube with Motion Blur, 8 samples, Bf=0.5.

Frame 1 of moving cube with Motion Blur, 8 samples, Bf=0.5. on the other hand shows the rendering of frame 1 when Motion Blur is set and 8 'intermediate' frames are computed. Bf is set to 0.5; this means that the 8 'intermediate' frames are computed on a 0.5 frame period starting from frame 1. This is very evident since the whole 'blurriness' of the cube occurs on half a unit before and half a unit after the main cube body.

Frame 1 of moving cube with Motion Blur, 8 samples, Bf=1.0. and *Frame 1 of moving cube with Motion Blur, 8 samples, Bf=3.0.* show the effect of increasing Bf values. A value greater than 1 implies a very 'slow' camera shutter.



Frame 1 of moving cube with Motion Blur, 8 samples, Bf=1.0.



Frame 1 of moving cube with Motion Blur, 8 samples, Bf=3.0.

Better results than those shown can be obtained by setting 11 or 16 samples rather than 8, but, of course, since as many *separate* renders as samples are needed a Motion Blur render takes that many times more than a non-Motion Blur one.

Hints

If Motion Blur is active, even if nothing is moving on the scene, Blender actually 'jitters' the camera a little between an 'intermediate' frame and the next. This implies that, even if OSA is off, the resulting images have nice Anti-Aliasing. An MBLUR obtained Anti-Aliasing is comparable to an OSA Anti-Aliasing of the same level, but generally slower.

This is interesting since, for very complex scenes where a level 16 OSA does not give satisfactory results, better results can be obtained using *both* OSA and MBlur. This way you have as many samples per frame as you have 'intermediate' frames, effectively giving oversampling at levels 25,64,121,256 if 5,8,11,16 samples are chosen, respectively.

Stamp

Blender 2.46+
Stamps the render with key date/time and other info. See Reference Manual for more info.

Oversampling

Mode: All Modes

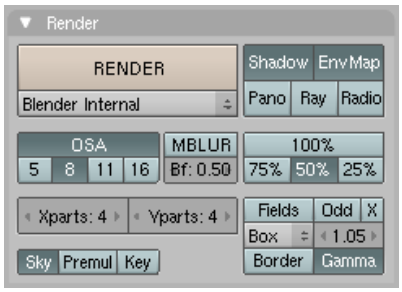
Panel: Render Context → Render

Hotkey: F10

Description

A computer generated image is made up of pixels, these pixels can of course only be a single colour. In the rendering process the rendering engine must therefore assign a single colour to each pixel on the basis of what object is shown in that pixel. This often leads to poor results, especially at sharp boundaries, or where thin lines are present, and it is particularly evident for oblique lines. To overcome this problem, which is known as *Aliasing*, it is possible to resort to an Anti-Aliasing technique. Basically, each pixel is 'oversampled', by rendering it as if it were 5 pixels or more, and assigning an 'average' colour to the rendered pixel. The buttons to control Anti-Aliasing, or OverSAmple (OSA), are below the rendering button in the Render Panel (*Render Panel*.).

Options



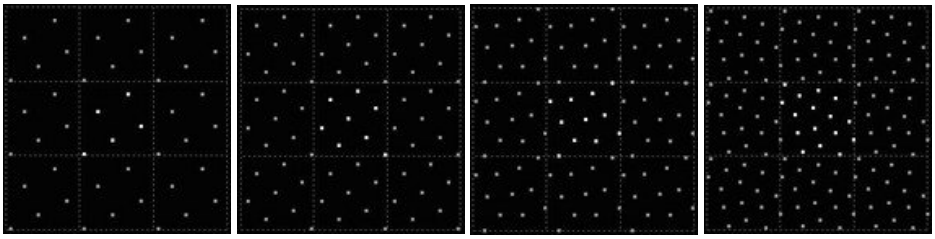
Render Panel.

OSA
: Enables oversampling

5 / 8 / 11 / 16
: The number of samples to use. The values of OSA (5, 8, 11, 16) are pre-set numbers in specific sample patterns; a higher value produces better edges, but slows down the rendering. By default, we use in Blender a fixed "Distributed Jitter" table. The samples within a pixel are distributed and jittered in a way that guarantees two characteristics:

1. Each sample has equal distances to its neighbour samples
2. The samples cover all sub-pixel positions equally, both horizontally and vertically

The images below show Blender sample patterns for 5, 8, 11 and 16 samples. To show that the distribution is equalized over multiple pixels, the neighbour pixel patterns were drawn as well. Note that each pixel has an identical pattern.



5 samples 8 samples 11 samples 16 samples

Filtering

When the samples have been rendered, we've got color and alpha information available per sample. It then is important to define how much each sample contributes to a pixel. The simplest method is to average all samples and make that the pixel color. This is called using a "Box Filter". The disadvantage of this method is that it doesn't take into account that some samples are very close to the edge of a pixel, and therefore could influence the color of the neighbour pixel(s) as well.

Filter menu

The filter type to use to 'average' the samples:

Box
: The original filter used in Blender, relatively low quality. For the Box Filter, you can see that only the samples within the pixel itself are added to the pixel's color. For the other filters, the formula ensures that a certain amount of the sample color gets distributed over the other pixels as well.

Tent
: A simplistic filter that gives sharp results

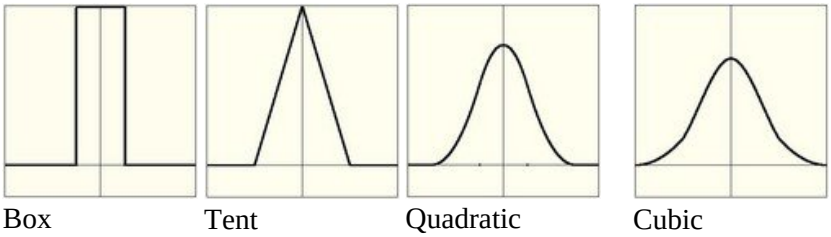
Quad
: A Quadratic curve

Cubic
: A Cubic curve

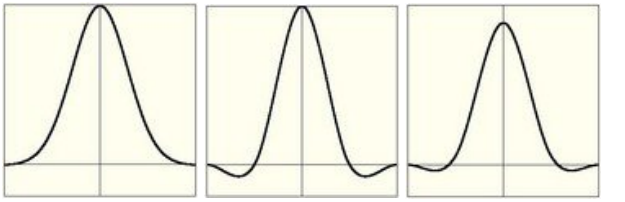
Gauss
: Gaussian distribution, the most blurry

CatRom
: Catmull-Rom filter, gives the most sharpening

Mitch
: Mitchell-Netravali, a good allround filter that gives reasonable sharpness



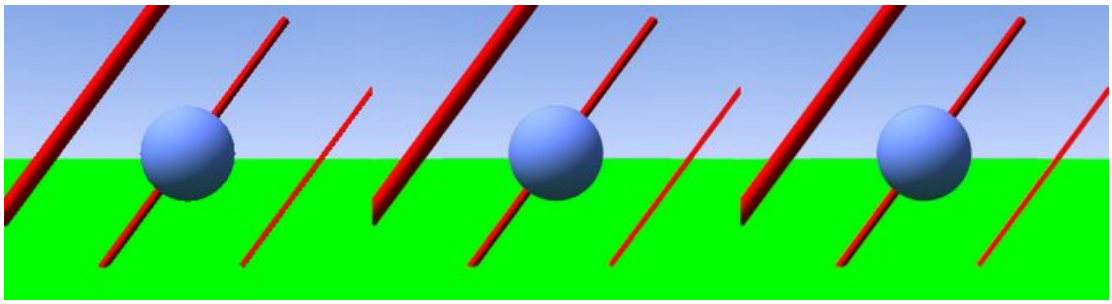
Box Tent Quadratic Cubic



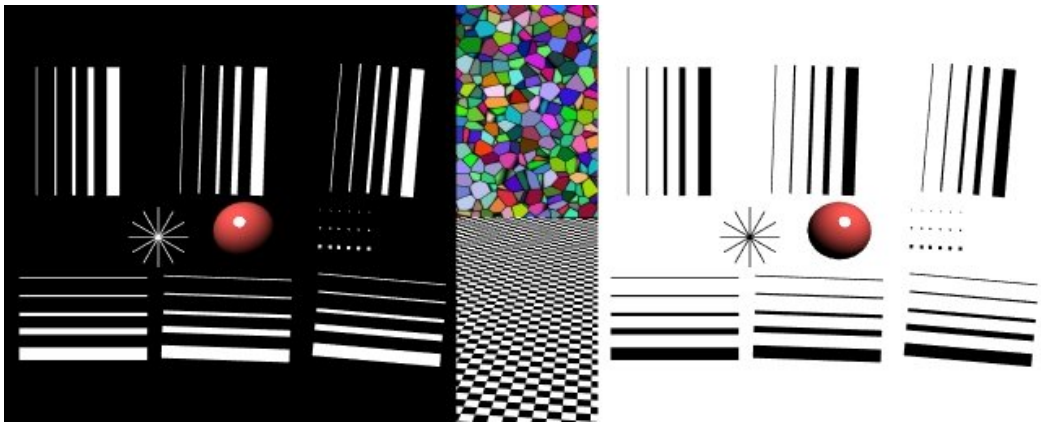
Gaussian Catmull-Rom Mitchell-Netravali

Making the filter size value smaller will squeeze the samples more into the center, and blur the image more. A larger filter size make the result sharper. Notice that the last two filters also have a negative part, this will give an extra sharpening result.

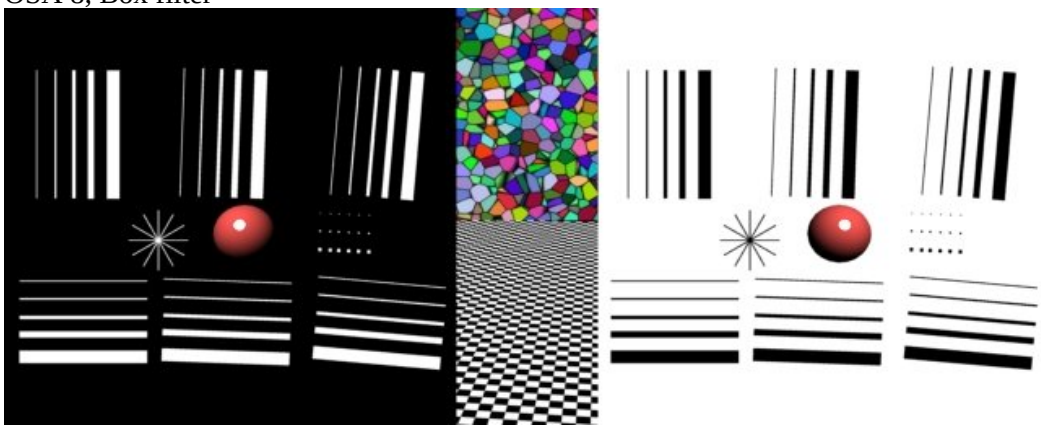
Examples



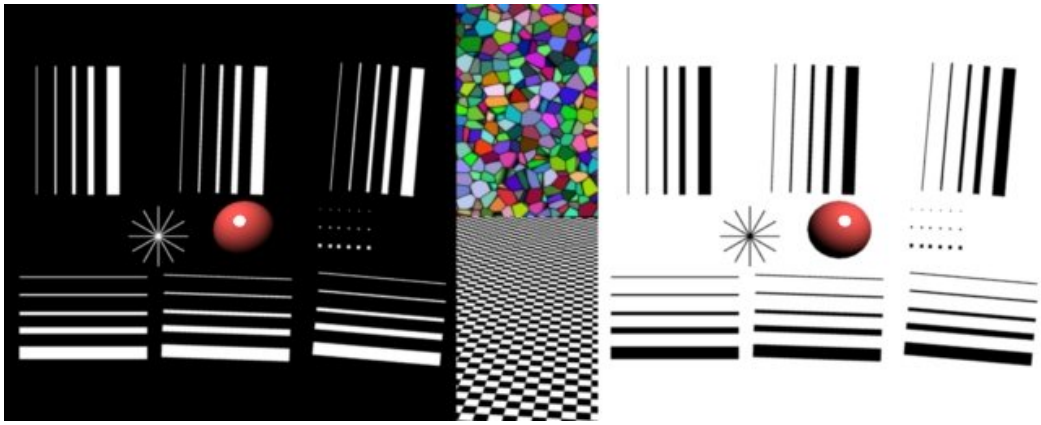
Rendering without OSA (left) with OSA=5 (center) and OSA=8 (right).



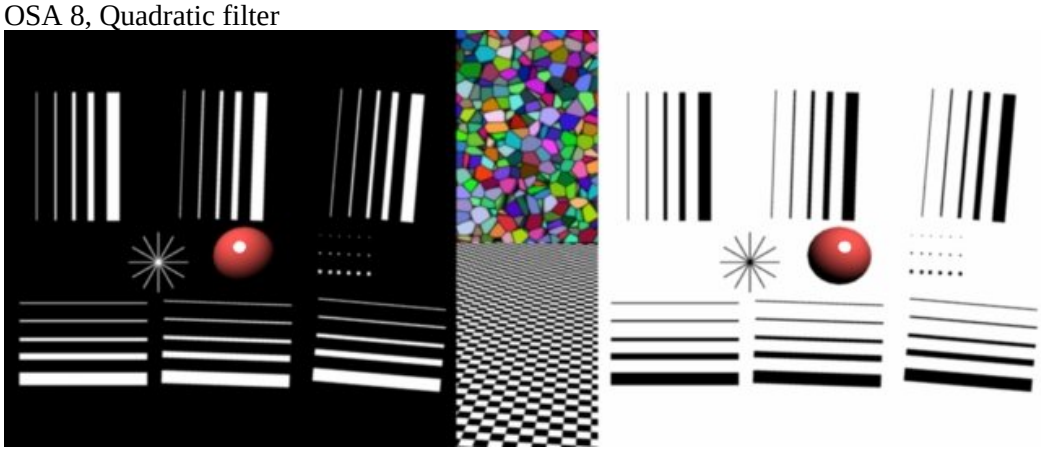
OSA 8, Box filter



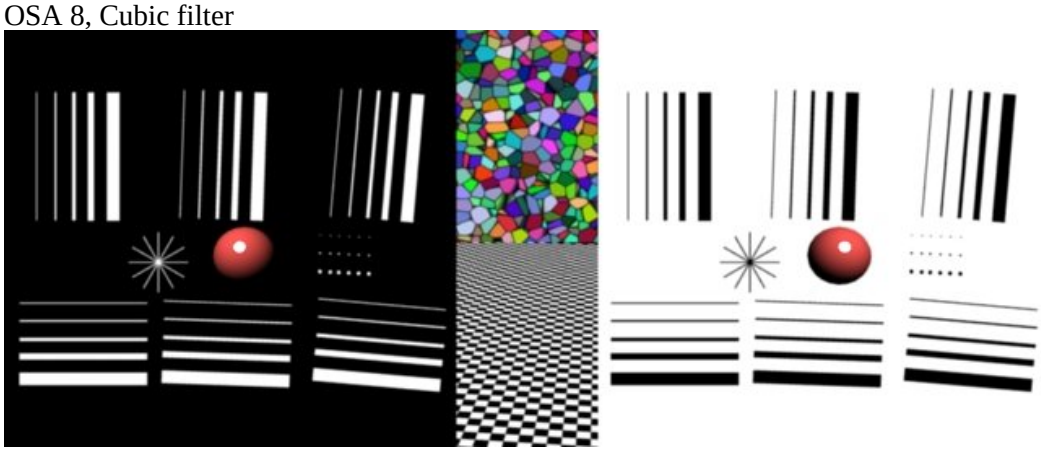
OSA 8, Tent filter



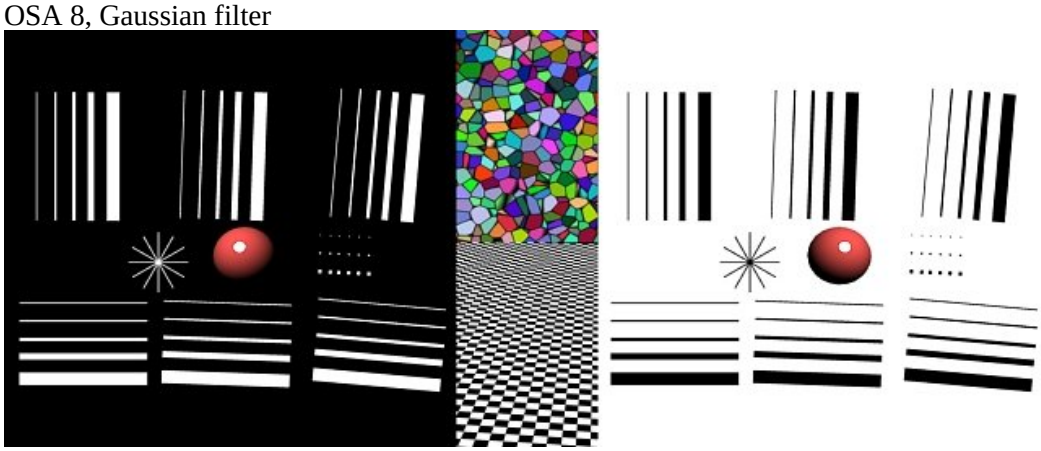
OSA 8, Quadratic filter



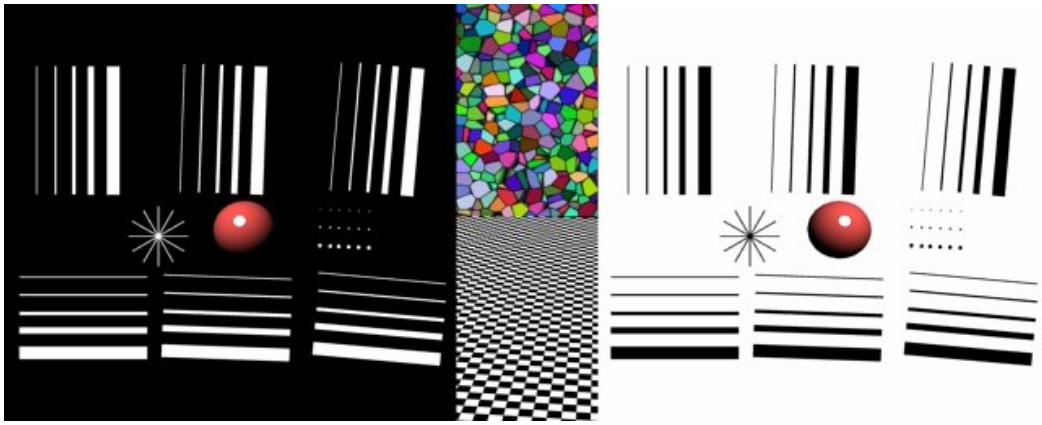
OSA 8, Cubic filter



OSA 8, Gaussian filter



OSA 8, Catmull-Rom filter



OSA 8, Mitchell-Netravali filter

Blender produces all the information needed to render a scene. While it has its own internal rendering engine, you can export or link to external renderers for image computaiton. Some external renderers include:

- [Yafray]
- [Yaf(a)ray]
- [Indigo]
- [Kerkythea]
- [PovRay]
- [MentalRay]
- [Sunflow]
- [Luxrenderer]

Of these, Blender is tightly integrated with Yafray, and an overview and quick start guide is provided here. For more detailed information, consult the *[Yafray web site].

Rendering with Yafray

Mode: Any Mode

Panel: Render Context → Render / Render Context → Yafray

Hotkey: F10

Description

Yafray, as the lengthened version of its name (Yet Another Free RAYtracer) suggests, is a free, XML speaking, cross platform raytracer developed by the Yafray team.. It works with many 3D modelling applications (with Wings and Aztec serving as examples), but the focus on this document shall fall upon it's use with Blender.

Yafray is currently available (under the LGPL license) for Windows, Linux (via source code compilation, or .deb or .rpm installation), Mac OSX, and Mac Intel; and installation packages, as well as Yafray's source code, can be downloaded here.

Options

All post-2.34 Blender releases have the option to call Yafray in the place of Blender's internal renderer, assuming it's installed. This can be done by switching to the "Render buttons" panel (F10) and selecting Yafray as the rendering engine.

Render Pipeline
When Yafray is used, it is inserted into the pipeline before any compositor or sequencer actions, because it is the renderer, and the compositor and sequencer work on rendered images. The image data given to Yafray is the scene objects, materials, lights, etc. Yafray does not know nor care about render layers and cannot feed Blender's node compositor or sequencer effects, since it takes a completely different approach and cannot produce the different render layers that the Blender internal renderer can. Yafray render frames based on Blender scene data.

To use Yafray with Blender's compositor, render the image using Yafray, and then use the image input node to get that image into the compositor where it can be post-pro. You can then feed that to Sequencer via the Scene strip and Do Composite enabled. To feed the Sequencer directly from Yafray's output, use the Image strip after Yafray has completed the render.

Two other panels should appear once Yafray's selected from this menu, which serve to supply a number of Yafray's options to you (other options are available exclusively as XML code).



Enabling Yafray

XML
: This button, if pressed, will export your scene to a .xml file in your system's 'tmp' directory before Yafray renders it. Useful if you wish to make modifications to the XML file, or render the scene from a command line interface. Should you wish, however, to view Yafray's progress during a render in Blender's render window, it's better to unclick this button.

AutoAA
: This options allows you to toggle between manual and automatic control of anti-aliasing options in the scene. Anti-aliasing is similar to Blender internals OSA, which in effect dictates the accuracy of the edges in the render.

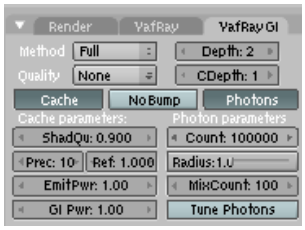
 In cases where you may need to manually control the anti-aliasing options (which becomes necessary if you wish to make use of Yafray's depth-of-field option), it's useful to remember that increasing the amount of samples per pass will increase the accuracy of the edges in the final render; decreasing the amount of samples per pass will, as you'd expect, decrease the accuracy, causing edges in the scene to seem rough, and jagged.

Proc.
: This option allows you to select the number of processors Yafray's allowed to make use of. For those of us who aren't lucky enough to have multiple processors, it's best to leave this option as it's default.

Gam.
: This option allows for manual correction of gamma values in the scene. The default (1) turns this option off.

Exp.
: This option allows for manual adjustment of exposure levels in the scene. A more indepth explanation of this option will come later.

Global Illumination



Global illumination settings
The next tab along, titled "Yafray GI", provides a selection of methods with which Yafray's able to light a scene. Methods available are:

Full
: This method works well for most scenes, most notably indoor scenes, where the use of photons becomes appropriate.

Skydome
: This method is more suited to outdoor scenes.

Cache
: Clicking the **cache** button speeds up rendering by allowing Yafray to be more selective it's distribution of samples. When this button's depressed, Yafray renders a pre-pass to determine the most suitable allocation of samples, before rendering the image itself, increasing the efficiency of the render.

The cache button then reveals three more options.

ShadQu
: This option allows for greater control over the quality of shadows. By increasing this option from its default (0.900), you also increase the number of samples taken in shadowed areas, which in turn not only increases the quality of shadows in the scene, but also increases render times.

Prec
: This option sets the maximum number of pixels per-square without samples. By decreasing this option from its default (10), you increase the number of samples taken in the scene. Decreasing this option also increases render times.

Ref
: This option allows the user to specify the threshold to refine shadows. By decreasing this option from its default (1.000), you invite Yafray to increase the number of passes taken to distribute samples in shadowed areas, thereby increasing the quality of the shadows in the scene, and increasing render times.

Examples

Starting with the default Blender set up, enable Yafray in the Rendering Buttons panel (F10), deselect the XML option in the "yafray" tab, and select the "full" method from the "yafray GI" tab, and set the quality to "low". Then click "Render" (F12).

Console output

Provided the evironment allows it, Yafray should output information to the console window (in Windows, Blender opens along side a console window by default. In GNU/Linux, however, to view the console output, you'll need to start Blender from the console - Usually by typing "blender" into a terminal emulator window).

If you switch to the console after the render's completed, you should (provided the "cache" option's enabled) notice something similar to this:

Console output

Launching 1 threads

Fake pass: [##############]
534 samples taken

Render pass: [##############]

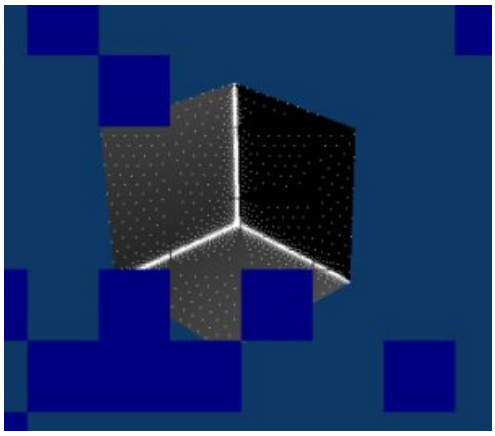
render finished

Output description

The render's split up into two seperate passes. The first, "fake" pass is made as a direct result of the "cache" option being enabled, and it's purpose is to determine the best distribution of samples in the scene (without the cache option enabled, the samples are distributed evenly among the scene). The number of samples is then output onto the next line.

The next pass is the "real" render pass, where Yafray renders the image based on the sample map created in the previous pass.

Render window output

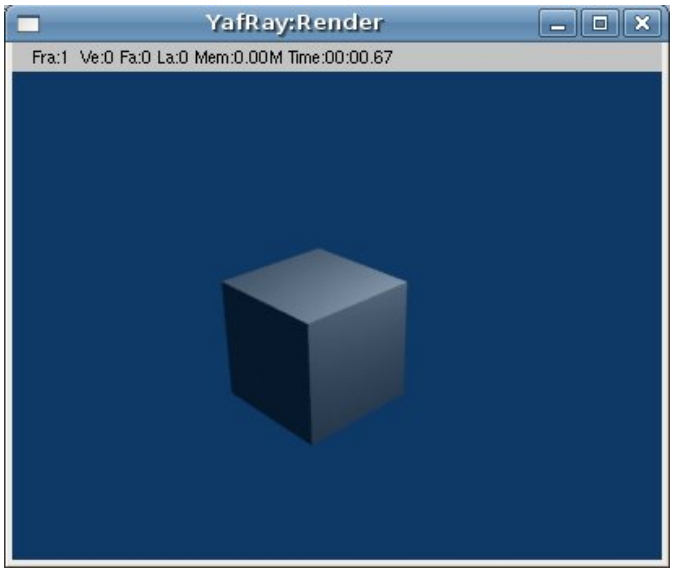


Greater samples in shadowed areas
Now we'll look at the Yafray's output to the render window, during the render.

Provided the XML option is turned off, Yafray will continually update it's visual output to the render window - Much like Blender does. The image to the right was captured during the "fake" pass stage of the render, and the white dots represent the allocation of samples in the scene. Notice how the samples are only placed on areas of the scene that are directly affected by light, meaning that, in the demonstration image, only the parts of the scene with a surface are considered.

This also means that in shadowed areas of the scene, the number of samples is greater.

You can notice that the density of white dots which, as I pointed out earlier, represent the number of samples per pixel in that area of the image, is greater in areas that are likely to be shadowed (in this case, I deleted the vertex of the cube closest to the camera, revealing inside edges, which aren't as exposed to the light).



Basic Yafray render

The rendered image

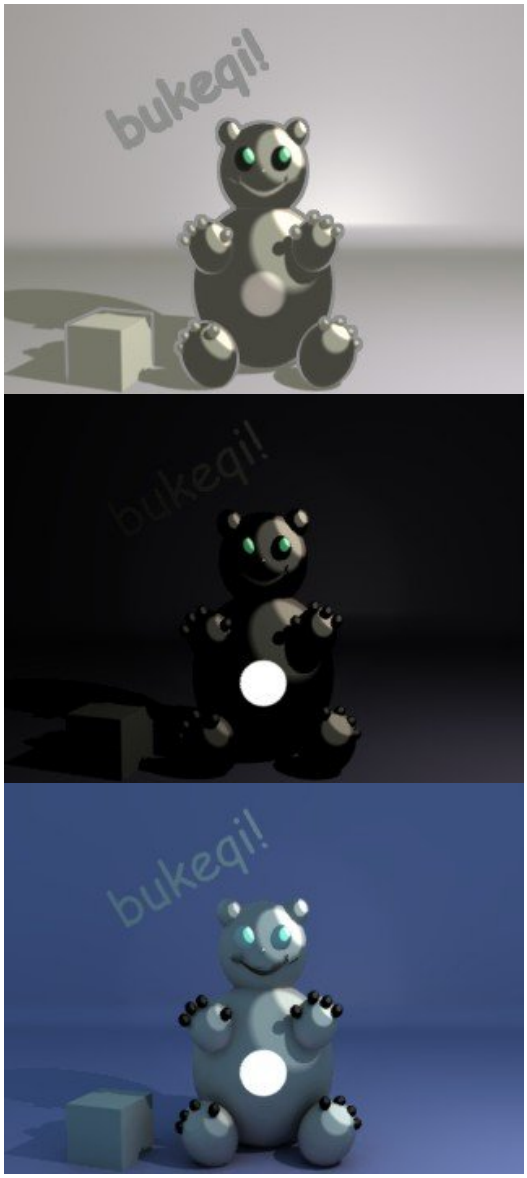
You'll notice how the cube, despite Blender's default grey material being applied, has been coloured blue. This is because the Full method is affected by the "world" colour of the scene, which, again as Blender's default, is blue. To change this, switch to the "shading" panel (F5), and select the little world icon. To have materials show properly, set the world shader to white.



Selecting the world shader

Notes

Amount of Light



YafRay deals with light completely differently than the Blender Internal Renderer, and apparently light intensity needs to be pumped by large amounts for YafRay. the images reflect a Blender Internal, a Yafray render without Global Illumination (GI), and one with Full GI. As you can see, results vary widely based on the illumination method chosen.

A solution is to use very large Area lamps (Square, 100 Size but Samples at only 4, Energy 10) for softer shadows, in combination with a Sun lamp at much lower Energy value (less than 1.0) if you want a distinct shadow edge. Sun lamps seem to provide much greater intensity than Area lamps in YafRay but the shadow edges are quite harsh.

Try using the Skydome setting for the YafRay GI because with Full GI you may get weird blotchy artifacts that no one seems to know how to remedy, but may be related to the scale of my Blender scene, which is 1BU = 1cm, with a figure built to life-size. You'll be doing something like this as well if you build a scale model to match camera perspectives.

Blender World parameters may include a small AO setting which YafRay does seem to take into account, so you might try adding some in your scene. Also be aware that the World Sky colors (Ho & Ze) are treated as a "hemi" light source, and will color your scene accordingly when using Skydome -- play with these RGB values to perhaps boost the overall lighting intensity by "filling in" with GI. In the pics below, the World lighting settings were doubled for the render on the right.

Everything seems to need to be boosted for YafRay -- some Materials look very dull unless you "double-up" some of the components (such as by using an image texture twice with "Add"), and the RGB & Shader tab settings are very different from what you would use with the Internal renderer.

You can also adjust the EmitPwr and Exp settings in the YafRay renderer tabs to compensate for the lighting differences. It gets to be quite a juggling act. The plus side is that you are able to get lighting of a much richer character for a scene, so it can be worth the trouble.

SkyDome



Effect of Blender "World" RGB settings on YafRay Skydome renderings

Various coloring effects based on World settings
Using the Blender Internal (BI) renderer, the only way to get the world Horizon, Zenith, or Textured color to affect the material color is to use Ambient Occlusion set to Sky Color or Sky Texture; otherwise (without AO) it only affects the color of the background. The only variable to directly affect the final object coloration in Blender Internal is the color of Ambient light, and then each material can receive a specified amount of that ambient light (by default 50%). The color of the ambient light in BI cannot be varied over the height of the image and is applied uniformly to the subject. Ambient Occlusion, based on the settings, affects the color of the model based on its geometry.

In Yafray, however, a key difference is that the color of all of these matter, as shown in the example. The example has the same material (the skin and hair) rendered using different **Horizon and Zenith** colors. Each of these, in effect, change the ambient light cast onto the subject. If the Zenith was darker, as is usually the reality, the tops of the model would be darker than the the lower portions. Using the color of the sky and horizon to affect the lighting of subjects lends a much more realistic blending of a subject into the environment, leading to more photorealistic results.

To achieve the same effect in Blender, you can use Ambient Occlusion, or light your subject with Hemisphere lamps which are the same color as your sky zenith and horizon.

Baking, in general, is the act of pre-computing something in order to speed up some other process later down the line. Rendering from scratch takes a lot of time depending on the options you choose. Therefore, Blender allows you to "bake" some parts of the render ahead of time, for select objects. Then, when you press Render, the entire scene is rendered much faster, since the colors of those objects do not have to be recomputed.

Render Baking

Mode: All Modes except Sculpt

Panel: Scene (F10) Render Context → Bake panel

Hotkey: CtrlAltB

Menu: Render → Bake Render Meshes

Description



The Bake tab in the Render buttons panel.
Render baking creates 2D bitmap images of a mesh object's rendered surface. These images can be re-mapped onto the object using the object's UV coordinates. Baking is done for each individual mesh, and can only be done if that mesh has been UV-unwrapped. While it takes time to set up and perform, it saves render time. If you are rendering a long animation, the time spent baking can be much less than time spent rendering out each frame of a long animation.

Use Render Bake in intensive light/shadow solutions, such as AO or soft shadows from area lights. If you bake AO for the main objects, you will not have to enable it for the full render, saving render time.

Use Full Render or Textures to create an image texture; baked procedural textures can be used as a starting point for further texture painting. Use Normals to make a low-resolution mesh look like a high-resolution mesh. To do that, UV-unwrap a high-resolution, finely sculpted mesh and bake its normals. Save that normal map, and Map To (material settings) the UV of a similarly unwrapped low-resolution mesh. The low-resolution mesh will look just like the high-resolution, but will have much fewer faces/polygons.

Advantages

- Can significantly reduce render times
- Texture painting made easier
- Reduced polygon count
- Repeated renders are made faster, multiplying the time savings

Disadvantages

- Object must be UV-unwrapped.
- If shadows are baked, lights and object cannot move with respect to each other.
- Large textures (eg 4096x4096) can be memory intensive, and be just as slow as the rendered solution.
- Human (labor) time must be spent unwrapping and baking and saving files and applying the textures to a channel.

You apply the baked Normal image as a Normal Map using a Texture channel.

Options

Full Render
: Bakes all materials, textures, and lighting except specularity and SSS.

Ambient Occlusion
: Bakes ambient occlusion as specified in the World panels (F8). Ignores all lights in the scene.

Normals
: Bakes tangent and camera-space normals (amongst many others) to an RGB image.

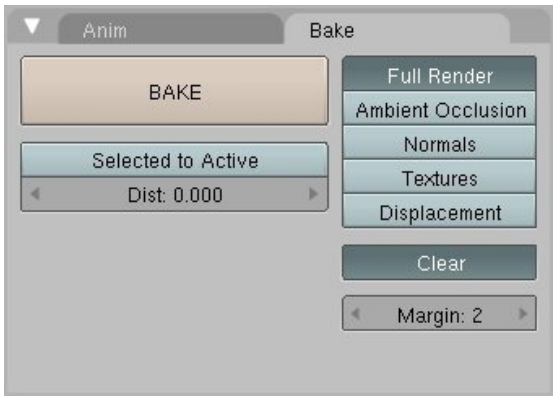
Textures
: Bakes colors of materials and textures only, without shading.

Clear
: If selected, clears the image to selected background color (default is black) before baking render.

Margin
: Baked result is extended this many pixels beyond the border of each UV "island," to soften seams in the texture.

Mesh Must be Visible in Render
If a mesh not is visible in regular render, for example because it is disabled for rendering in the Outliner or has the DupliVerts setting enabled, it cannot be baked to.

Options under development (after 2.45)



The Bake tab in the Render buttons panel.***Development version***
Read more here: Changes since 2.45 -> Baking Selected to Active

Selected to Active toggle button
: Information from other object can be baked onto the active object, with the Selected to Active option.

Distance
: The Distance parameter controls how far a point on another object can be away from the point on the active object. Only needed for Selected to Active.

 A typical use case is to make a detailed, high poly object, and then bake it's normals onto an object with a low polygon count. The resulting normal map can then be applied to make the low poly object look more detailed.

Displacement Map (new option for render-bake type)
: Similar to baking normal maps, displacement maps can also be baked from a high-res object to an unwrapped low-res object, using the Selected to Active option.

When using this in conjunction with a subsurf and displacement modifier within Blender, it's necessary to temporarily add a heavy subsurf modifier to the 'low res' model before baking. This means that if you then use a displacement modifier on top of the subsurf, the displacement will be correct, since it's stored as a relative difference to the subsurfed geometry, rather than the original base mesh (which can get distorted significantly by a subsurf). The higher the render level subsurf while baking, the more accurate the displacements will be. This technique may also be useful when saving the displacement map out for use in external renderers.

Normals can now be baked in different spaces:

- Camera space - The already existing method.
- World space - Normals in world coordinates, dependent on object transformation and deformation.
- Object space - Normals in object coordinates, independent of object transformation, but dependent on deformation.
- Tangent space - Normals in tangent space coordinates, independent of object transformation and deformation. This is the new default, and the right choice in most cases, since then the normal map can be used for animated objects too.

For materials the same spaces can be chosen as well, in the image texture options, next to the existing Normal Map setting. For correct results, the setting here should match the setting used for baking. Note that this replaces the NMap TS setting in material options, which is automatically converted to the Tangent space option in the texture options.

Workflow

Windows Users do AO first
If you are running Blender on a Windows operating system, you may have to first bake Ambient Occlusion before baking any other option. If you do not bake AO first, you may get the error message **"No Image to Bake To"** and will not be able to bake anything for that mesh and will have to restart Blender.

1. In a 3D View window, select a mesh and enter UV/Face Select mode
2. Unwrap the mesh object
3. In a UV/Image Editor window, either create a new Image or open an existing Image. If your 3D view is in Textured display mode, you should now see the image mapped to your mesh. Ensure that all faces are selected.
4. With your mouse cursor in 3D View, press CtrlAltB to popup the menu of available baking choices. Alternatively, access the Bake panel in the Buttons window, Scene (F10) context, Render sub-context.
5. Bake your desired type of image: Full Render, Ambient Occlusion, Normals, or shadeless Textures.
6. After computation, Blender replaces the image with the Baked image.
7. Save the image in the UV/Image Editor window via Image->Save
8. Refine the images using the process described below, or embellish with Texture Paint or an external image editor.
9. Apply the image to the mesh as a UV texture. For displacement and normal maps, refer to Bump and Normal Maps. For full and texture bakes, refer to Textures. In all cases you want to Map Input to the UV Map that was used when baking.

Render baking and additional options are available in the Render buttons context, Bake tab. The saved image can now be assigned as an image texture using UV coordinates in Map Input tab of Materials context (F5). In a multiple UV setup, images can be baked from one (inactive) UV set to another (active) set. This can be useful for seam removal, or for game engines that require particular UV setups.

Refining Baked Images

Mode: UV/Image Editor

Hotkey: AltS

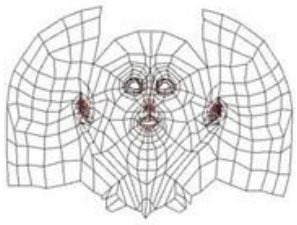
Menu: Image->Save

After baking, be sure to save the generated image. Images may also be refined in the UV/Image Editor using Texture Paint.

Examples

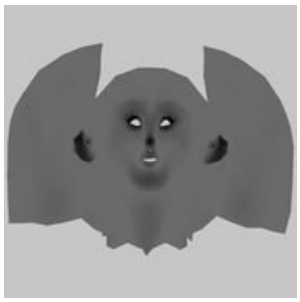


"Adrianna" model by Mr_Bomb; unwrap and textures by BgDM



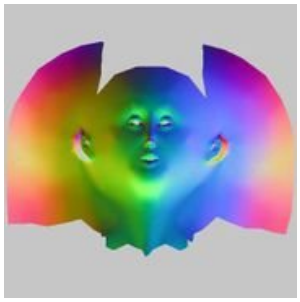
UV Layout for Adrianna model

These examples use a model of a human head, and its UV Layout, shown to the left. The mesh was seamed along the top scalp line and down back of the head, and unwrapped using LSCM projection. It has seven layered image and procedural textures, and is lit by a hemisphere of spots as well as several accent lights.



Ambient occlusion bake

Ambient Occlusion option: Generates a bitmap of the "unwrapped" ambient occlusion result, ignoring all lighting. AO maps are often multiplied with color channels to simulate self-shadowing, soft lighting, or dirt. Note: if the texture shown here were simply multiplied by the diffuse texture, it would darken it by half; see below for a compositing noodle to brighten and blur this image.



Normals bake (camera space)

A Normals bake generates a camera-space normals RGB map. The orientations of the mesh's vertex normals relative to the Camera's X, Y, and Z axes are recorded as different levels of Red, Green, and Blue respectively. In this example, the camera was aimed at the right side of the head model, so the right-side faces are largely blue since they faced the camera more or less straight-on (normals parallel to camera's Z). Camera-space maps have limited utility; this map would allow you to create relief on a lower-poly model, but the illusion would be accurate only if viewed from the right side of the head. The solution to this is a "Tangent space" normal map which is independent of camera position and object deformation. Blender can bake many different types of maps, and to see them, please go to Doc:Manual/Textures/Maps/Bump and Normal Maps.



Textures only bake

The Textures option renders all materials and textures devoid of lighting, AO and shadows (similar to a "shadeless" render). This example baked image comprises the color and bump textures, two Cloud textures, and a subtle linear blend texture. A potential use of such a map might be to multiply it with a baked AO map in the compositor.



Full render bake

Full Render bake takes into account lighting/shadow, materials, vertex colors, bump, and AO if applicable. Full render baking does not include specularity and subsurface scattering, as these are view-dependent parameters. Neither does this mode support halo materials. Full renders can be useful as time-savers for renders of static scenes or objects; for example, soft shadows of a tree or piece of furniture cast by an area light (slow to render) can be baked and "painted" into the scene (very fast render), and will remain convincing as long as the object or light do not move.



Margin option "bleeds" pixels beyond edges of UV map

The Margin option adds the specified number pixels beyond the edge of the UV map. Can be useful for minimizing seams when the image is mapped back onto the mesh. If no value is specified, 1 pixel is added by default.

Applying Baked Images

When you baked the mesh, you were in UV Face Select Mode and the image is created. Save or Pack the image. Then, depending on the Option chosen, you map the rendered texture to the mesh using the following method:

Full Render

In the object's material settings, enable TexFace, which tells Blender not to compute the material

color based on the procedural material settings, but instead use the UV Texture image mapped to the mesh according to the UV Layout.
Alternatively, create a Texture channel for the Material. In the MapInput panel, enable UV and enter the name of the UV Texture ("UVTex" by default). In the MapTo panel, enable "Col" for Color.
Activate the Texture buttons subcontext, choose Image as the type of texture, and Load the image by selecting it from the selector list by the button "Load New".
With either mapping method, TexFace or Texture Channel, since the Full Render also takes into account the lighting, you should also enable Shadeless in the Material panel. When rendering the composite scene, Blender does not waste time recomputing the lighting effects on Shadeless materials, saving a lot of render time.

Ambient
: Add a texture channel, and in the MapInput panel, map the texture to UV as described above. In the MapTo panel, enable "Col" but change the mix method to "Multiply". Blender will now multiply the above texture color channels by the AO mask, darkening the colors in the cracks and crevices.

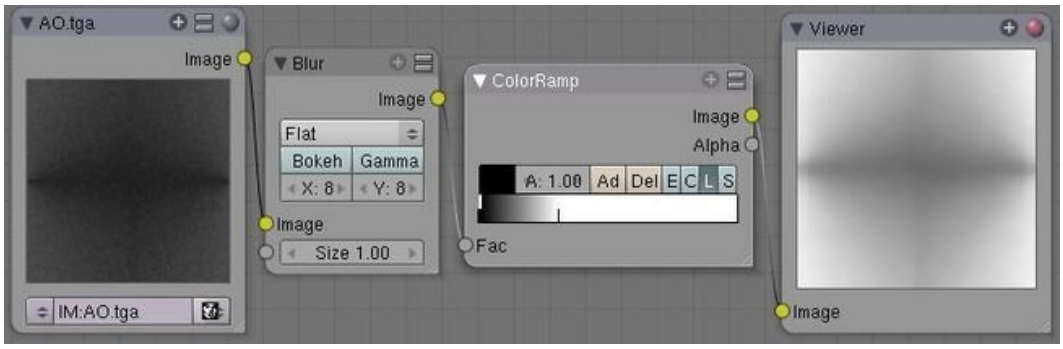
Normal
: Add a texture channel and load the baked normals image. Make sure to activate "Normal Map" in the Map Image panel. In the Materials contents Map Input panel, enable UV; in the Map To panel, enable Nor and turn off "Col". For camera mapped normals, do not use NMap TS in the shader; this option is for tangent space maps only

Texture

The Texture bake includes all materials, but no shading information.

Refining the Ambient Occlusion Image

Often, the rendered AO image will be dark and grainy. Keep in mind this image is multiplied by the base material, so if it is dark, it will darken the whole mesh. If it is grainy, it will make your mesh look speckled or dirty. If so, try increasing the number of *Samples* and disabling *Noise*.



Suggested compositing noodle for AO images

We recommend that you save the AO bake as an image file named "<object>-AO.<ext>", where <object> is the name of the selected object, and <ext> is the appropriate format extension ("tga" for targa format, "jpg" for JPEG, etc). Note that if you have extensions enabled, Blender will add the correct <ext> for you.

As a last resort, although not technically precise, you can blur the image to smooth it out, and increase the dynamic range of the image. To do so, switch to Composite Nodes and construct the noodle shown above. Read in that AO file using the Image input node. The Blur node will remove the graininess, and the adjust the ColorRamp node so that the brightest parts of the image are white (a value of 1.0).

Using a UV/Image Editor window, select "Viewer Node" as the display, and you can then save the image back out; we suggest using something like "<object>-AOK.<ext>". Now, use the AOK file as the multiplier against Color, and it will delicately and smoothly shade the mesh for ambient occlusion.

Applying High-Res Normal to a Low-Res Mesh

One of the tricks of the trade, especially in gaming, is to make a low-poly model look better by modeling two versions of the same thing; one low-poly, and another high-poly. (Blender's "multi-resolution mesh" feature is specifically designed for this.) Then, you bake the normal map of the high-poly model, which gives you a texture image that shows all the shadows and creases and detail. Then, apply those images to the low-poly model, and you save yourself a lot of time that you would otherwise have to spend hand-painting the low-poly UV Texture. For detailed explanations on normal maps, please visit Doc:Manual/Textures/Maps/Bump and Normal Maps.

Thanks to the option which allows baking from selected to Active, you do not need to unwrap the high-poly mesh. Here are the steps:

1. Make a copy of your object.
2. Unwrap the copy at a low resolution.
3. Add a new image to the unwrapped mesh using the UV / Image window.
4. Make sure the High-poly and Low-poly objects share the exact same location in Blender (having them on separate layers can help. You can work on one layer at a time and then select both layers when you bake.)
5. Select the High-Poly Object, Shift Select the Low-Poly object. This leaves both selected with the Low-poly being the Active selected object.
6. In the Render panel, select the Bake tab.
7. Choose the map you want to Bake (Normal or Displacement).
8. Choose the "Selected to Active" button.
9. Click Bake.

Hints

Currently (v2.45), only mesh objects can be render-baked.

Baked images are not automatically saved. Save from the UV editor window with Alt-S or Image menu -> Save.

Baking can be interrupted with the Esc key.

If you move lights, you should re-bake any Full Render textures. If you change the mesh geometry, you should re-bake all of the options.

If you change any object's material color or any settings, you have to re-bake that object's Texture or Full Render for those changes to take effect.

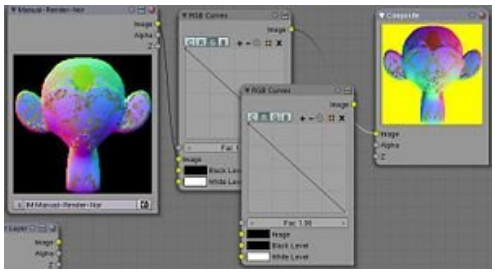
Procedural textures (e.g. Clouds) can be baked with the Full Render or Textures options, providing a basis for further detailing. Use Shadeless materials if necessary.

If you get the error "No images found to bake to", check the following:

1. The object is selected
2. The object is **visible** (check the Logic panel)

3. Ensure there is a new image (with UV test grid) loaded for baking.
4. If Step 1 does not work although object meets all criteria in that Step, switch to Object mode, select the object, and try to bake.
5. Try Image->New again in the UV editor and bake to AO before anything else.

Blender to xNormal/Maya Normal



Converting Blender Normal to xNormal
Blender's normal map baking tool is different from other applications, such as xNormal and Maya. Both of these are widely used in the games industry and are already production proven (meaning the way they bake their normal maps is more commonly used and can be regarded as a standard). To convert a Blender normal map to an xNormal, you need to invert the Red and Green channels. This is easily accomplished using the noodle shown to the right. Use the Image input node to read in the baked normal from Blender. Thread the Image socket to an RGB Curves node. In the Red and Green channels, drag the left endpoint up, and the right endpoint down, so that the curve runs diagonally down as shown. Thread the Image output to the Composite output, enable Do Composite, and save the image.

Introduction

In some situations we want to increase the render speed, access blender remotely to render something or build scripts that use blender command line.

One advantage of using command line is that we don't need the X server (in case of Linux) and as a consequence we can render remotely by SSH or telnet.

Note! Arguments are executed in the order they are given!

```
blender -b file.blend -a -x 1 -o //render
```

...Wont work, since the output and extension is set after blender is told to render.

Always position **-f** or **-a** as the last arguments.

Description

Usage: blender [-b <dir><file> *[-o <dir><file>*][-F <format>][-x *[0|1]*][-t <threads>][-S <name>][-f <frame>][-s <frame> -e <frame> -a]] [[-P <scriptname> [-- <parameter>]]

Render options:

```
 -b <dir><file> Render <file> that is inside the <dir> without load the UI
   -P <filespec>Run the specified Python script (filename or Blender Text)
   -S <name>    Set scene <name>

   -f <frame>   Set frame <frame> to render and save it (Don't use together with -a)
   -j <number>  Render every x frames (jump frames by this number)

   [-s <frame>] [-e <frame>] -a
                Set start frame (-s), end frame (-e) or both.
                The order is important and it's possible to use only the -s or the -e
                alone.

   -o <dir><file>
                Set the render path and file name.
                Use // as <dir> to use the path render relative to the blend file.
                Use # in the filename to be replaced with the frame number
                  eg: blender -b foobar.blend -o //render_# -F PNG -x 1

   -F <format>  Set the render format, Valid options are..
                TGA IRIS HAMX FTYPE JPEG MOVIE IRIZ RAWTGA
                AVIRAW AVIJPEG PNG BMP FRAMESERVER
      (formats that can be compiled into blender, not available on all systems)
       HDR TIFF EXR MPEG AVICODEC QUICKTIME CINEON DPX

   -x [0|1]     Set option to add the file extension to the end of the file, 0 means
                no and 1 means yes.

   -t <threads> Use amount of <threads> for rendering
```

Animation options: (Used when pressing the play button, blender behaves like a movie player)

```
 -a <file(s)>  Playback <file(s)> (Don't use together with -b!)
   -p <sx><sy> Open with lower left corner at <sx>,<sy> (Doesn't work on win)
   -m          Read from disk (Don't buffer)
```

Window options:

```
 -w            Force opening with borders (default)
 -W            Force opening without borders (Linux/Unix Only)
 -p <sx> <sy> <w> <h>  Open with lower left corner at <sx>, <sy>
                       and width and height <w>, <h>
```

Game Engine specific options:

```
 -g fixedtime          Run on 50 hertz without dropping frames
 -g vertexarrays       Use Vertex Arrays for rendering (usually faster)
 -g noaudio            No audio in Game Engine
 -g nomipmap           No Texture Mipmapping
 -g linearmipmap       Linear Texture Mipmapping instead of Nearest (default)
```

Misc options:

```
 -d            Turn debugging on
```

```
-noaudio      Disable audio on systems that support audio
-h            Print this help text
-y            Disable script links, use -Y to find out why its -y
-P <filename> Run the given Python script (filename or Blender Text)
-R            Register .blend extension
-v            Print Blender version and exit
```

Examples

Render a picture

```
# blender -b file.blend -o //file -F JPEG -x 1 -f 1
```

- **-b file.blend** : File .blend to render
- **-o //file** : Directory + Target image file
- **-F JPEG** : JPEG image format
- **-x 1** : Adds an extension .jpg to the file name
- **-f 1** : Render frame 1

Render a movie

```
# blender -b file.blend -x 1 -o //file -F MOVIE -s 003 -e 005 -a
```

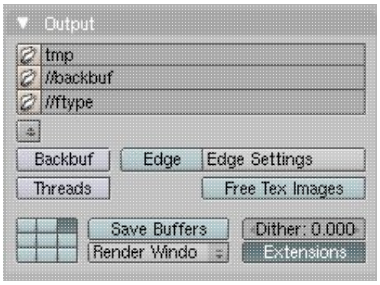
- **-b file.blend** : File .blend to render
- **-x** : Adds an extension .avi to the movie
- **-o //images/file** : Directory + Target image file
- **-F MOVIE** : This saves a .AVI movie with low compression
- **-s 003 -e 005 -a** : Set start frame to 003 and end frame to 005. ***Important:*** *You can use -s or -e, but if they're not in order, they'll not work!*

Output Panel

Mode: All Modes

Panel: Render Context → Render

Hotkey: F10



Render Output buttons.
This panel provides many options for rendering, increasing and optimizing your render and output speed, and the location for displaying and saving your render output. The options on this panel control where and how the results of a render are handled.

File Locations

At the top are three filespec fields:

- Animation Output Directory and filename prefix (default: //tmp)
- Background image (default: //backbuf)
- FType image format descriptor file (default: //ftype)

By default, each frame of an animation is saved in the directory where the file was loaded from (the "//" part of the path specification) and given a filename that begins with "tmp". Change this or any field by â § ShiftLMB clicking in the name field and entering a new name. If you use the // and do not save a new .blend file somewhere, Blender assumes the // to refer to the Blender install folder.

A background image, such as a studio curtain, a watermark, or any image may be used as a background. You usually want to set the window background to this picture, and when you render, use the background image in the render output.

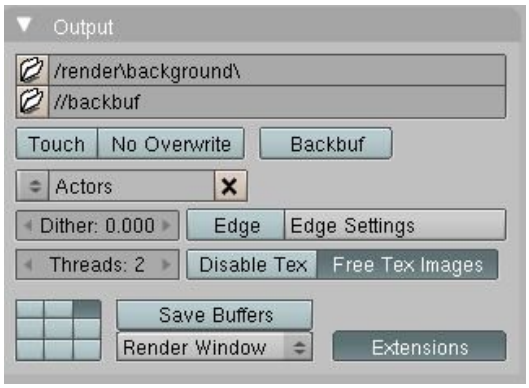
Ftype uses an "Ftype" file, to indicate that this file serves as an example for the type of graphics format in which Blender must save images. This method allows you to process 'color map' formats. The color map data is read from the file and used to convert the available 24 or 32 bit graphics. If the option "RGBA" is specified, standard colour number '0' is used as the transparent colour. Blender reads and writes (Amiga) IFF, Targa, (SGI) Iris and CDinteractive (CDi) RLE colormap formats.

Directory Browser
Clicking the folder icon to the left of the field turns a Blender window pane into a File Browser window. Using this window is very handy for scrolling through your hard disk and selecting a file or directory.

PathSpecs
The path specification for the location can include a normal or mapped drive letter (e.g. "F:"), a breadcrumb notation (e.g. "./" and "../" and "//" (file location). Forward slashes (Unix-style) or backslashes (Windows-style) are acceptable on either platform. If omitted, the file is saved in the Blender Installation

directory.

Scene Linking



Background scene linking in Actors Scene
The up-down button just below the ftype field lists the scenes in the file and allows you to link in another scene. When you select a scene, you will note that all the objects in that scene are shown in your 3D windows. They have a special outline and you cannot select them, but they are shown and will participate in the render. The linked scene can in turn link to a third scene, and so on infinitum; you can form a chain, linking scenes together for an ordered set of scenes to all be rendered together. This feature almost makes scenes akin to layers.

Use this feature for rendering complicated scenes, yet retaining the ability to quickly work on and test render only one major set of objects at a time by un-linking. The render is relative to the current scene's camera, so as long as all objects are positioned properly relative to 3D Space

This feature is useful for teamwork. For example, one team can be working on props and background, while another team works on characters. The backgrounds team can, at any time, link to the actor's scene, and make sure that the actors are not walking through tables and walls, etc.

Other Buttons and Options

Backbuf
: Clicking the Backbuff (off by default) makes blender use the image file specified above in the backbuf field as a background image.

Edge
: Edge button enables cartoon-style outlining of edges of objects. When clicked and enabled, The edge settings dynamic menu allows you to set the Edge Intensity (Eint) of the edge and its color via the RGB sliders (default:black).

Threads:
: Threads enables multi-threaded rendering; great for multi-core and multi-processor systems. On an multi-processor PC, this allows the rendering work (a Thread) to be divided among each processor. If you have a dual-core PC (Intel or AMD), set this to 2 (and increase Xparts and Yparts)

Disable Tex
: Using textures, either procedural or image, takes a tremendous amount of processing power. To speed up draft renders, enable this button

Free Tex Images

To save memory, enable the Free Tex Images so that texture images are cleaned out of memory.

Render Location button grid

The 3x3 Grid specifies where (generally) the render window should pop up.

Save Buffers

Save Buffers saves the intermediate layers and render pipeline images to disk as an OpenEXR file, in order to free up physical memory. Use this if you are running out of memory; there will be more hard disk cache used and hence rendering may take little longer because of I/O increases, but you will be able to render HUGE images on your PC.



Error: Sorry...exr tile saving only allowed with equally sized parts:

The XParts and YParts must divide the image up into square patches with respect to pixel size. For example, a 1920x1080 HD image is in a 16:9 ratio; 120 times 16:9 gives 1920x1080. So, an XParts of 16 and YParts of 9 would divy up the render into 144 patches that are 120 pixels square.

Other combinations for XParts:YParts for any sized image in a 16:9 aspect ratio are 64:36 (2304 patches), 80:45 (3600 patches), and 96:54 for over 5000 patches. Using XParts of 80 and YParts of 45, you can render a 3 million polygon image at 50% of HD using full textures and very complicated lighting (ED Scene 7 Shot 4) on a PC with 2G RAM.

Render Display

Render Display is a dynamic menu that lets you select where the image should be displayed

- ◊ Render Window (default) - In a pop-up window.
- ◊ Full Screen - Full-screen on your monitor.
- ◊ Image Editor - Piped directly into Blender's built-in Image Editor.

When rendering to an Image Editor window or a popup Render Window, the grid of 9 buttons allows you to select the location of the window, either the window frame that will be used to become the Image Editor window, or the general location of the popup Render Window when it pops up.

Dither: setting

Dither numeric field (0 is off) specifies the amount of noise to inject into the picture.

Extensions toggle button

Extensions tells blender to add a file extension onto the output file based on the format of the picture. For example, frame 35 of a series of JPEG-encoded pictures would be named "tmp0035.jpg". You almost always want Blender to add the filetype extension so that the image is "recognized" by your operating system.

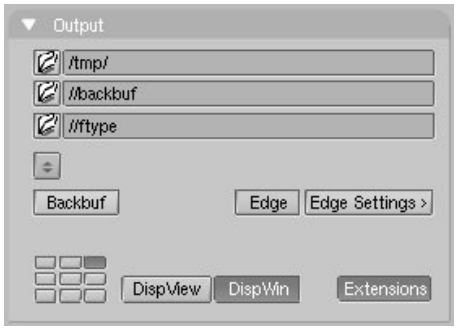
Edge (Toon) Rendering

Description



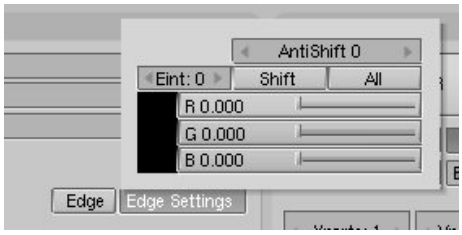
A scene with Toon materials.
Blender's toon shaders, can give your rendering a comic-book-like or manga-like appearance, affecting the shades of colours, as you may be able to appreciate in *A scene with Toon materials.*. The effect is not perfect since real comics and manga also usually have china ink outlines. Blender can add this feature as a post-processing operation.

Options



Toon edge buttons *(F10)*.

Edge
: This makes Blender search for edges in your rendering and add an 'outline' to them.



Toon edge settings *(F10)*.
Before repeating the rendering it is necessary to set some parameters. The Edge Settings opens a window to set these (*Toon edge settings (F10).*). Edge Settings

Colour / R/G/B
: The colour of the rendered edges (black by default). Use the sliders or click on the swatch to see the

color picker

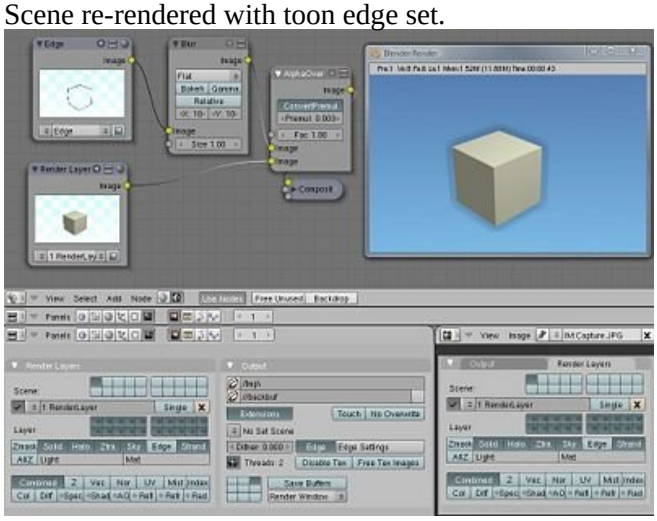
Eint

The edge's intensity from 0 to 255. 10 gives outline of object against the background, whereas higher settings start to pick up forward or leading edges based on the contrast in the image caused by the geometry of the object and not specular spots caused by lighting. At maximum intensity, Edge will even faintly display geometry subsurf edge lines in areas of imperfect smoothing.

Examples



Scene re-rendered with toon edge set.



It is possible to separate out the edge layer using a renderlayer dedicated to that purpose. The alpha channel is 0 where there is no edge, and 1 where the edge is. By separating out the edge layer, you can blur it, change its color, mask it, etc. The image to the right shows how to do this. I created an Edge renderlayer that only has the Sky and Edge layers (I included sky so that we get the world color later on in the composite output). The other renderlayer omits the Edge layer, so it returns just the normal image. On the output panel I enabled Edge with a width of 10 in black. I run that layer through a blur node. Using the Alphaover node, I then composite the cube on top of the blurred edge. The result gives a soft-shadow kind of effect. Note that Premultiply is set because the Edge image already has an alpha of 1.0 set.

Output Formats

Mode: All Modes

Panel: Render Context → Format

Hotkey: F10

Description

The output format for either single image renders (**RENDER** F12 + F3 save image) or for Animations **ANIM** CtrlF12 is selected using the Format Panel. From here you can select many image or animation formats (*Image and animations formats.*).

There are many image formats out there for many different uses. A format stores an image in a *lossless* or lossy format; with lossy formats you suffer some image degredation but save disk space because the image is saved using fewer bytes. A lossless format preserves the image exactly, pixel for pixel. You can break formats down into *static* images and movie *clips*. Within either category there are standards (static formats and clip codecs) which may be proprietary standards (developed and controlled by one company), or open standards (which are community or consortium-controlled). Open standards generally outlive any one particular company and will always be royalty-free and freely obtained by the viewer. Proprietary formats may only work with a specific video card, or while the codec may be free, the viewer may cost.

Format Panel

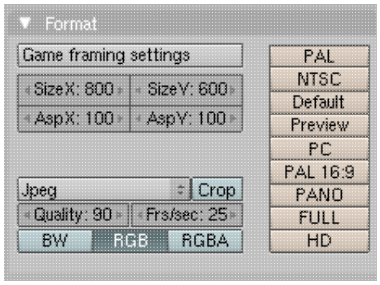


Image Format options.
By default the dimensions SizeX and SizeY are 800Ã 600 and can be changed as for any Num Button by shift-clicking into the field. These buttons control the overall size of the image. Just below are two more settings, AspX and AspY which control the packing of the pixels along the respective axis. Together, the four buttons below define the viewing size and aspect ratio of the image. Aspect is the ratio between the X and Y dimensions of the pixel of the image. By default it is 1:1 since computer screen pixels are square. If television shorts are being made, and since TV pixels are not square, you want to change this aspect ratio to match the destination video standard: PAL for Europe, and NTSC for the Americas.

Presets

To make life easier the rightmost block of buttons provide some common presets:

- PAL - 720Ã 576 pixels at 54:51 aspect ratio.
- NTSC - 720Ã 480 pixels at 10:11 aspect ratio.
- Default - Same as PAL, but with full TV options, as explained in the following sections.

- Preview - 640Ã 512 at 1:1 aspect ratio. This setting automatically scales down the image by 50%, to effectively produce a 320Ã 256 image.
- PC - 640Ã 480 at 1:1 aspect ratio.
- PAL 16:9 - 720Ã 576 at 64:45 aspect ratio, for 16:9 wide screen TV renderings.
- PANO - Standard panoramic settings 576Ã 176 at 115:100 aspect ratio. More about 'panoramic' renderings in the pertinent section.
- FULL - 1280Ã 1024 at 1:1 aspect ratio.
- HD - Maximus Pixelus. With a 16:9 aspect ratio and checking in at a full 1920Ã 1080 resolution for every frame, this is the ultimate format, with over 2 million pixels to compute for every frame. Bassam says to "Grab a cup of coffee while you wait, but the results are spectacular." You can also set format to 1280x720 30 fps to get HD 720p flavor. Check the Even Interlacing button for the "i" variety of 1080/720, thus giving support for all HighDef varieties.

These are just the presets; you can set any resolution you wish, subject to your PC's memory restrictions; see the Render page for ideas and techniques and tools for enabling huge render outputs.

Options

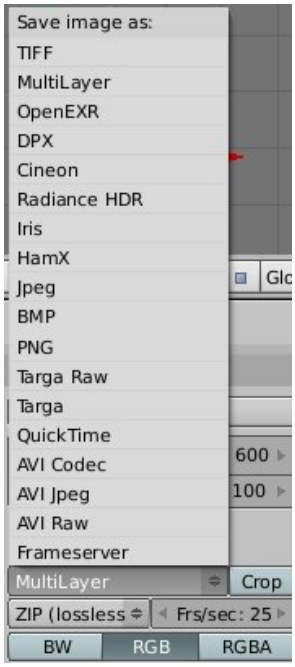


Image and animations formats.
Blender supports a wide mix of image formats. Some formats are produced by internal Blender code. Others (Tiff, for example) may require a dynamic load library (such as libtiff.dll) in your Blender installation folder. In alphabetical order they are (**bold** indicates a movie clip format):

- **AVI Raw** - Audio-Video Interlaced (AVI) uncompressed frames.
- **AVI Jpeg** - AVI but with Jpeg compression. Lossy, smaller files but not as small as you can get with a Codec compression algorithm. Jpeg compression is also the one used in the DV format used in the digitals camcorders.
- **AVI Codec** - AVI codec compression. Available codecs are operating system dependant. When an AVI codec is initially chosen, the codec dialog is automatically launched. The codec can be changed directly using the Set Codec button which appears (*AVI Codec settings.*).
- BMP Bit-Mapped Paint lossless format used by early paint programs.
- Cineon - format produced by a Kodak Cineon camera and used in high-end graphics software and more directed toward digital film.

- DPX - Digital Moving-Picture eXchange format; an open professional format (close to Cineon) that also contains metainformation about the picture; 16-bit uncompressed bitmap (huge file size). Used in preservation.
- Frameserver - Blender puts out frames upon request as part of a render farm. The port number is specified in the OpenGL User Preferences panel.
- HamX - Blender's own self-developed 8 bits RLE (Run Length Encoded bitmap) format; it creates extremely compact files that can be displayed quickly. To be used only for previsualization of animations (Play button).
- Iris - the standard Silicon Graphics Inc (SGI) format used on those spanking Unix OS machines.
- Jpeg - Joint Picture Expert Group (name of the consortium which defined it), an open format that supports very good compression with little loss of quality. Only saves RGB value. Re-saving images results in more and more compression and loss of quality.
- MultiLayer - an OpenEXR format that supports storing multiple layers of images together in one file. Each layer stores a renderpass, such as shadow, specularity, color, etc. You can specify the encoding used to save the MulitLayer file using the codec selector (ZIP (lossless) is shown and used by default).
- OpenEXR - an open and non-proprietary extended and highly dynamic range (HDR) image format, saving both Alpha and Z-depth buffer information.

 ◊ Enable the *Half* button to use 16-bit format; otherwise 32-bit floating point precision color depth will be used
 ◊ Enable the *Zbuf* button to save the Z-buffer (distance from camera) info
 ◊ Choose a compression/decompression *CODEC* (ZIP by default) to save disk space.
 ◊ Enable the *RGBA* button to save the Alpha channel.
 ◊ Because OpenEXR is so new and previews are generally not supported by Operating Systems, enable *Preview* to save a JPG image along with the EXR image so you can quickly and easily see what the basic image looks like.

- PNG - Portable Network Graphics, a standard meant to replace old GIF inasmuch as it is lossless, but supports full true colour images. Supports Alpha channel.

 Enable the RGBA button to save the Alpha channel.

- Radiance HDR - a High Dynamic Range image format that can store images with big changes in lighting and color.
- TARGA and Targa raw - Truevision Advanced Raster Graphics Adapter is a simple raster graphics format established in 1984 and used by the original IBM PC's. Supports Alpha Channel.

 Enable the RGBA button to save the Alpha channel.

- TIFF - Often used for teletype and facsimile (FAX) images
- **QuickTime** - Apple's Quicktime .mov file. The Quicktime codec dialog is displayed when this codec is installed and this format is initially chosen.

Reads GIF if QuickTime is Installed
Blender can read GIF files on Windows and Mac platforms with [QuickTime] installed. The GIF capabilities (as well as flattened PSD, flattened PDF on Mac, and others) come along with QuickTime.

Compression

Some formats can compress the image to use less disk space. This compression might be lossless (PNG, ...) or lossy (Jpeg, ...). Lossy formats don't store individual pixel information, thus reducing image quality. All the other formats are more or less equivalent, each having advantages and disadvantages. Make your compression selection using the button or field located beneath the format selector. For example, if Jpeg is selected, you

can specify a compression level (Quality:90 by default). Higher quality takes more disk space, but results in a better looking picture with less compression encoding artifacts.

The default image type is Targa, but, since the image is stored in a buffer and then saved, it is possible to change the image file type after the rendering and before saving using this menu. (**attention**: this is only valid for static images, not when rendering animations!).

Channels

Blender renders color (RGB) images, but Black and White (BW) and color with Alpha Channel (RGBA) are also possible. Beware, unless the Extensions button of the Output pannel is set, Blender does *not* automatically add the extension to files, hence any .tga or .png extension must be explicitly written in the File Save window.

If the option "RGBA" is specified, standard colour number '0' is used as the transparent colour. The format MUST support an alpha channel as part of its specifcation; for example, if you choose a Jpeg format (whose specifcation does not support alpha transparency), clicking RGBA will not magically add an alpha channel; it will be saved as RGB only.

OpenEXR and **Multilayer** formats are the only formats that store Z-depth buffer information. **Multilayer** is the only format that stores Render Layer and Render Passes as channels that can then be composited in post-production.

Blender reads and writes (Amiga) IFF, Targa, (SGI) Iris and CDinteractive (CDi) RLE colormap formats. Specify the colormap file in the Output FType field.

Frame Rate

For an animation the frame rate (*AVI Codec settings.*) which, by default, is 25 frames per second, the standard for PAL (European) television. Use 30 frames per second for USA television.

Codec



AVI Codec settings.
A Codec is a little routine that compresses the video so that it will fit on a DVD, or be able to be streamed out over the internet, over a cable, or just be a reasonable file size. Codecs compress the channels of a video down to save space and enable continuous playback. *Lossy* codecs make smaller files at the expense of image quality. Some codecs, like H.264, are great for larger images. Codecs are used to encode and decode the movie, and so must be present on both the encoding machine (Blender) and the target machine. The results of the encoding are store in a container file.

Blender knows two kinds of container files:

- Audio Video Interlace (a .avi extension) and
- QuickTime (a .mov extension).

When AVI Codec is selected, Blender will popup a little Codec selector window, listing the codecs that are registered on your machine. Each Codec has unique configuration settings. Consult the documentation on the codec (supplied by the company that wrote it) for more information.

When Quicktime is selected, the codecs on your machine will pop-up and allow you to pick which one you want to use. You may have to have purchased Quicktime Pro to use this.

There are dozens, if not hundreds, of codecs, including XviD, H.264, DivX, Microsoft, and so on. Each has advantages and disadvantages and compatibiliity with different players on different operating systems.

Most codecs can only compress the RGB or YUV color space, but some support the Alpha channel as well. Codecs that support RGBA include:

- animation (quicktime)
- PNG *TIFF *Pixlet - not lossless, and may be only available on Apple Mac.
- [Lagarith Lossless Video Codec]

More information on image formats can be found at:

- http://www.cineon.com/
- http://www.digitalpreservation.gov/

Interlacing

Interlacing is a way of providing a sort of motion blur with compression. Instead of capturing the full resolution image every so many times a second, half of the horizontal scan lines are captured twice the number of times a second. So, instead of displaying 1280x720 images 25 times a second, which is called HD 720p EU, you could display 1280x360 images 50 times a second, where the first frame is the even scan lines (horizontal rows 2, 4, 6, 8 ...) and the second frame is set 1/50th of a second later consisting of the odd scan lines (rows 1, 3, 5, 7,...). The net result is that the same number of pixels are displayed every second, but the interlaced variety will appear smoother, since the odd lines catch any movement that happened in between the even frames, and vice versa. Blender supports Even interlacing (described above, used for EU TV) and Odd interlacing, for US TV, where the first frame is scan lines (1, 3, 5, 7 ...) and the frame after is the even lines (2, 4, 6, 8, ...). Use the Even/Odd buttons for this purpose.

Frame Rate

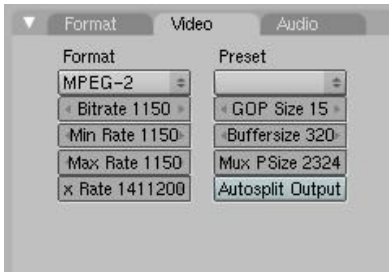
Additionally, Blender supports 50 or 60, or 24 or 30 frames per second. 50 and 25 fps are used for EU TV, and 60 and 30 are US, and 24 is film. Because of US Power Grid cycles, the actual frame rate is 29.97 fps. To accommodate this, Blender has a divider for the frame rate field; enter 30 fps and a divider of 1.001 to get exactly 29.97 fps.

FFMPEG

FFMPEG, short for Fast Forward Moving Pictures Expert Group, is a collection of free and open source software libraries that can record, convert and stream digital audio and video in numerous formats. It includes libavcodec, an audio/video codec library used by several other projects, and libavformat, an audio/video container mux and demux library.

When you select FFMPEG as your output format, two more tabs appear that allow you to select the video codec and the audio codec.

FFMPEG Video

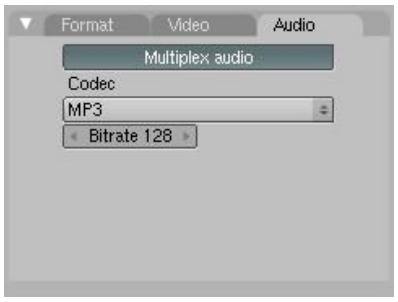


Here you choose which video codec you want to use, and compression settings. With all of these compression choices, there is a tradeoff between filesize, compatibility across platforms, and playback quality. You can use the presets, DV, SVCD, DVD, etc. which choose optimum settings for you for that type of output, or you can manually select the format (MPEG-1, MPEG-2, MPEG-4, AVI, Quicktime (if installed), DV, H.264, or Xvid (if installed). You must have the proper codec installed on your computer for Blender to be able to call it and use it to compress the video stream.

If your video is HUGE and exceeds 2Gig, enable Autospilt Output. The main control over output filesize is the GOP, or keyframe interlace. A higher number generally leads to a smaller file, but needs a higher-powered device to replay it.

Standards
Codecs cannot encode off-the-wall video sizes, so stick to the XY sizes used in the presets for standard TV sizes.

FFMPEG Audio



Audio is encoded using the codec you choose, as long as you enable Multiplex Audio. For each codec, you may be able to control the bitrate (quality) of the sound in the movie. This example shows MP3 encoding at 128kbps. Higher bitrates are bigger files that stream worse but sound better. Stick to powers of 2 for compatibility.

Free Advice

Choosing which format to use depends on what you are going to do with the image. If you are going to

- email it to your friends, use **JPG**
- combine it with other images in post processing and simple color/alpha composition, use **PNG**
- use nodes to simulate depth of field and blurring, use **EXR**

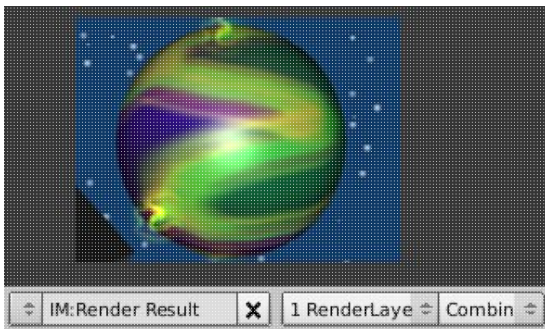
- composite using Render Passes, such as the Vector pass, use **Multilayer**.

If you are animating a movie and are not going to do any post-processing or special effects on it, use either **AVI-JPEG** or **AVI Codec** and choose the XviD open codec. If you want to output your movie with sound that you have loaded into the VSE, use **FFMPEG**.

If you are going to do post-processing on your movie, it is best to use a frame set rendered as **PNG** images; if you only want one file, then choose **AVI Raw**. While AVI Raw is huge, it preserves the exact quality of output for the post-processing. After post-processing (compositing and/or sequencing), you can compress it down. You don't want to post-process a compressed file, because the compression artifacts might throw off what you are trying to accomplish with the post-processing.

Note that rendering an animation long to calculate in a unique file (AVI or QuickTime) is more risky than in a set of static images: if a problem occurs while rendering, you have to re-render all from the beginning, while with static images, you can restart the rendering from the place (the frame) where the problem occurred!

Checking Render Results



When Blender renders an image, it saves it in an internal buffer. To see the latest render, change one of your panes to a UV/Image Editor. In the IM: Image selector, select Render Result, and the window will display the latest render. To the right of the Image selector, you can also select a RenderLayer to view, and, if you have render passes enabled, the results of a specific pass.

Rendering Animations

Mode: All Modes

Panel: Render Context → Anim

Hotkey: F10

Description

While rendering stills will allow you to view and save the image from the render buffer when it's complete, animations are a series of images, or frames, and are automatically saved directly out to disk after being rendered.

When using Blender, the idea is to use the Compositor to do green screen masking, matting, color correction, DOF, and so on to the images. That result is then fed to the Sequencer where the strips are cut and mixed and a final overlay is done.

Workflow

Generally, you do a lot of intermediate renders of different frames in your animation to check for timing, lighting, placement, materials, and so on. At some point, you are ready to make a final render of the complete animation for publication.

There are two approaches you can use when making a movie, or animation, with or without sound. The approach you should use depends on the amount of CPU time you will need to render the movie. You can render a "typical" frame at the desired resolution, and then multiply by the number of frames that will ultimately go into the movie, to arrive at an total render time. If the total render time is an hour or more, you want to use the "Frame Sequence" approach. For example, if you are rendering a one-minute video clip for film, there will be (60 seconds per minute) * (24 frames per second) or 1440 frames per minute. If each frame takes 30 seconds to render, then you will be able to render two frames per minute, or need 720 minutes (12 hours) of render time.

Rendering takes all available CPU time; you should render overnight, when the computer is not needed, or set Blender to a low priority while rendering, and work on other things (be careful with the RAM space!).

The **Direct Approach** is where your set your output format to an AVI or MOV format, and click ANIM to render your scene directly out to a movie file. Blender creates one file that holds all the frames of your animation. You can then use Blender's VSE to add an audio track to the animation and render out to an FFMPEG format to complete your movie.

The **Frame Sequence** approach is where you set your output format to a still format (such as JPG, PNG or MultiLayer), and click ANIM to render your scene out to a set of images, where each image is the frame in the sequence. Blender creates a file for each frame of the animation. You can then use Blender's compositor to perform any frame manipulation (post processing). You can then use Blender's VSE to load that final image sequence, add an audio track to the animation, and render out to an FFMPEG format to complete your movie. The Frame Sequence approach is a little more complicated and takes more disk space, but gives you more flexibility.

Here are some guidelines to help you choose an approach.

Direct Approach
- short segments with total render time < 1 hour
- stable power supply
- computer not needed for other uses

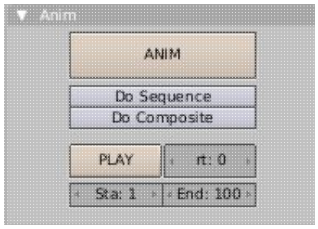
Frame Sequence Approach
- total render time > 1 hour
- post-production work needed
 - Color/lighting adjustment
 - Green screen / matte replacement
 - Layering/compositing
 - Multiple formats and sizes of ultimate product
- intermediate frames/adjustments needed for compression/codec
- precise timing (e.g. lip-sync to audio track) needed in parts
- may need to interrupt rendering to use the computer, and want to be able to resume rendering where you left off.

Frame Sequence Workflow

1. . First prepare your animation.
2. . In the Scene panels setup your animation to be rendered out as a image, generally using a format that does not compromise any quality (I prefer PNG or MultiLayer because of their lossless nature).
3. . Choose the output path and filespec in the Output render panel, for example "//\render\my-anim-"
4. . Confirm the range of your animation frame Start and End
5. . Save your .blend file.
6. . Press the big *ANIM* (animate) button. Sleep or get a life while you wait for your computer to do its thing.
7. . Once the animation is finished, use your OS file explorer to navigate into the output folder (".\render in this example). You will see lots of images (.png or .exr, etc... depending on the format you chose to render) that have a sequence number attached to them ranging from 0000 to a max of 9999. These are your single frames.
8. . In Blender, now go into the video sequence editor.
9. . Choose *Add Image* from the add menu. Select all the frames from your output folder that you want to include in your animation (Press A to Select All easily). They will be added as a strip to the sequence editor.
10. . Now you can edit the strip and add effects or simply leave it like they are. You can add other strips, like an audio strip.
11. . Scrub through the animation, checking that you have included all the frames.
12. . In the Scene Render buttons, in the Anim panel, below the animation button, activate *Do Sequence*.
13. . In the Format panel, choose the container and codec you want (e.g. FFMPEG H.264) and configure it.
14. . Click the ANIM button and Blender will render out the sequence editor output into your movie.

Why go through all this hassle you may ask? Well, first of all, if you render out single frames you can stop the render at any time by pressing Esc in the render window. You will not lose the frames you have already rendered, since they have been written out to individual files. You can always adjust the range you want to continue from where you left off. You can edit the frames after wards and postprocess them. You can add neat effects in the sequence editor. You can render the same sequence into different resolutions (640x480, 320x240, etc) and use different codecs (to get different file sizes and quality) with almost no effort whatsoever.

Options



Animation rendering buttons.

ANIM
: Starts the animation rendering.

Do Sequence
: Renders the output of the sequence editor, instead of the view from the 3D scene's active camera. If the sequence contains scene strips, these will also be rendered as part of the pipeline. If Do Composite is also enabled, the Scene strip will be the output of the Compositor.

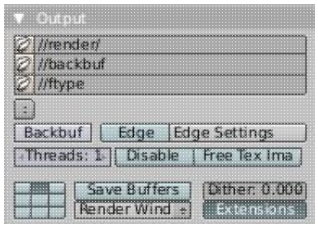
Do Composite
: Renders the output from the Compositing noodle, and then pumps all images through the Composite node map, displaying the image fed to the Composite Output node.

Play
: Opens a modal render window and plays the animation for you. You must close this window before returning to Blender via Esc or using the window control X. A nice feature of this window is that the frame number is shown in the header. If you LMB click in the window, the video will reset. Holding down your up-arrow will move the video forward in real-time; holding the down arrow reverses the video, thus allowing you to scrub through the video to identify exactly any errant frames.

Carried Away
If you render out to a frame sequence of 100 frames for example, and then change your Animation Sta: and End: to something smaller, say 1 to 30, Blender will play all the frames it can find in that sequence (all 100). If you re-render with changes to update that smaller subset, Blender will still play the new set plus the rest of the old set. The only solution is to modify the output file name before re-rendering!

Return Code
: Useful for debugging

Sta and End
: The start and end frame numbers to render the animation from and to. Frame numbers are inclusive. These are also set if you modify the Sta: and End: fields in the Timeline window.



Animation location and extensions.
By default the animation is rendered in the directory specified in the Output Panel (*Animation location and extensions.*). If an AVI format has been selected, then the name will be ####_####.avi where the '####' indicates the start and end frame of the animation, as 4 digit integers padded with zeros as necessary.

If an image format is chosen, on the other hand, a series of images named ####, ('####' being the pertinent frame number) is created in the directory.

Extensions
: Adds the correct file extensions per file type to the output files

Hints

Argh! My bratty sister turned off the PC right in the middle of rendering my movie!
: Unless your animation is really simple, and you expect it to render in half an hour or less, it is always a good idea to render the animation as separate image frames in a lossless format (TGA, PNG, BMP) rather than as a movie file from the beginning. This allows you an easy recovery if there is a problem fails and you have to re-start the rendering, since the frames you have already rendered will still be in the Output directory. Just change the STArt frame number to the frame number where you want to pick up from, and click ANIM again.

I only need to re-render a few frames in the middle
: It's also a good idea to render initially to a frame sequence, since if only a few frames have an error, you can make corrections and re-render just the affected frames. You can then make a movie out of the separate frames with Blender's sequence editor or with compositing nodes.

Only first frame renders, then Blender locks up
: If you click ANIM and only the first frame renders, be sure the output file is not locked by the media player. In general, check the console when rendering.

Unable to create Quicktime movie
: CreateMovieFile error: -47
: The Quicktime movie strip is in use (possibly in the VSE) and cannot be overwritten. If it is used in the VSE, delete the strip, or delete the file using your file explorer.

Preparing your work for video

Once you have mastered the trick of animation you will surely start to produce wonderful animations, encoded with your favourite codecs, and possibly you'll share them on the Internet with the rest of the community.

Sooner or later you will be struck by the desire of building an animation for Television, or maybe burning your own DVDs. To spare you some disappointment, here are some tips specifically targeted at Video preparation. The first and principal one is to remember the double dashed white lines in the camera view! If you render for PC then the whole rendered image, which lies within the *outer* dashed rectangle will be shown. For Television, some lines and some part of the lines will be lost due to the mechanics of the electron beam scanning in your TV's cathode ray tube. You are guaranteed that what is within the *inner* dashed rectangle in camera view will be visible on the screen. Everything within the two rectangles may or may not be visible, depending on the given TV set that your audience watches the video on.

The rendering size is strictly dictated by the TV standard. Blender has three pre-set settings for your convenience:

- PAL 720x576 pixels at 54:51 aspect ratio.
- NTSC 720x480 pixels at 10:11 aspect ratio.
- PAL 16:9 720x576 at 64:45 aspect ratio, for 16:9 widescreen TV renderings.
- HD 1920 x 1080 pixels at 1:1 aspect, that operates in a downscaled mode of 720 horizontal scan lines interlaced.

Aspect Ratio
TV screens do *not* have the square pixels which Computer monitors have; their pixels are somewhat rectangular, so it is necessary to generate *pre-distorted* images which will look bad on a computer but which will display nicely on a TV set.

If you render your animation at 1600x1200 resolution, and then burn a DVD, your image will not be clearer or crisper on the TV; in fact the DVD burning sofware will have had to downsize your images to fit the resolutions shown above, and you will have wasted about 4x disk space and render time.

Colour Saturation

Most video tapes and video signals are not based on the RGB model but on the YCrCb model: more precisely, the YUV in Europe (PAL), and the YIQ in the USA (NTSC), this latter being quite similar to the former. Hence some knowledge of this is necessary too.

The YCrCb model sends information as 'Luminance', or intensity (Y) and two 'Crominance' signals, red and blue (Cr and Cb). Actually a Black and White TV set shows only luminance, while colour TV sets reconstruct colour from Crominances (and from luminance). The contruction of the YCrCb values from the RGB ones takes two steps (the constants *in italics* depend on the system: PAL or NTSC):

- First, the Gamma correction (*g* varies: 2.2 for NTSC, 2.8 for PAL):
 - $R' = R^{1/g}$
 - $G' = G^{1/g}$
 - $B' = B^{1/g}$
- Then, the conversion itself:
 - $Y = 0.299R' + 0.587G' + 0.114B'$
 - $Cr = a_1(R' - Y) + b_1(B' - Y)$
 - $Cb = a_2(R' - Y) + b_2(B' - Y)$

Whereas a standard 24 bit RGB picture has 8 bits for each channel, to keep bandwidth down, and considering that the human eye is more sensitive to luminance than to chrominance, the luminance signal is sent with more bits than the two chrominance signals. This bit-expansion results in a smaller dynamic of colours, in Video, than that which you are used to on Monitors. You hence have to keep in mind not all colours can be correctly displayed. Rule of thumb is to keep the colours as 'greyish' or 'unsaturated' as possible, this can be roughly converted in keeping the dynamics of your colours within 80% of one another. In other words, the difference between the highest RGB value and the lowest RGB value should not exceed 0.8 ([0-1] range) or 200 ([0-255] range). This is not strict, something more than 0.8 is acceptable, but an RGB display with a color contrast that ranges from 0.0 to 1.0 will appear to be very ugly (over-saturated) on video, while appearing bright and dynamic on a computer monitor.

Rendering to fields

Mode: All Modes

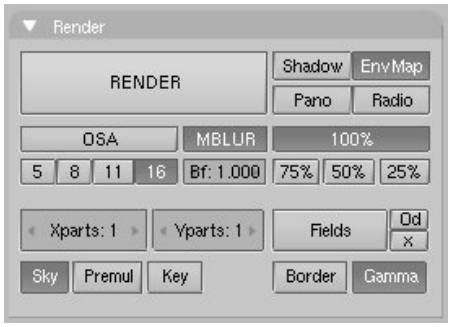
Panel: Render Context → Render

Hotkey: F10

Description

The TV standards prescribe that there should be 25 frames per second (PAL) or 30 frames per second (NTSC). Since the phosphorous of the screen does not maintain luminosity for very long, this could produce a noticeable flickering. To minimize this TVs do not represent frames as a Computer does ('progressive' mode), but rather represents half-frames, or *fields* at a double refresh rate, hence 50 half frames per second on PAL and 60 half frames per second on NTSC. This was originally bound to the frequency of power lines in Europe (50Hz) and the US (60Hz). In particular fields are "interlaced" in the sense that one field presents all the even lines of the complete frame and the subsequent field the odd ones. Since there is a non-negligible time difference between each field (1/50 or 1/60 of a second) merely rendering a frame the usual way and splitting it into two half frames does not work. A noticeable jitter of the edges of moving objects would be present.

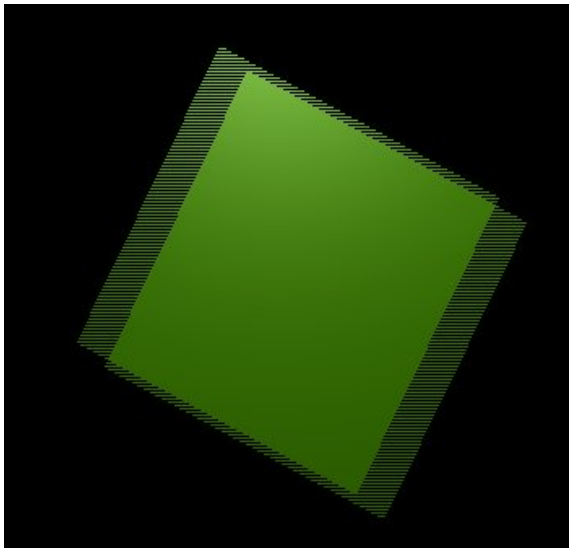
Options



Field Rendering setup.

Fields
: Enable field rendering. When the

Fields button in the Render Panel is pressed (*Field Rendering setup.*), Blender prepares each frame in two passes. On the first it renders only the even lines, then it *advances in time by half a time step* and renders all the odd lines.



Field Rendering result.
This produces odd results on a PC screen (*Field Rendering result.*) but will show correctly on a TV set.

Odd
: Forces the rendering of Odd fields first

X
: Disables the half-frame time step between fields (x).

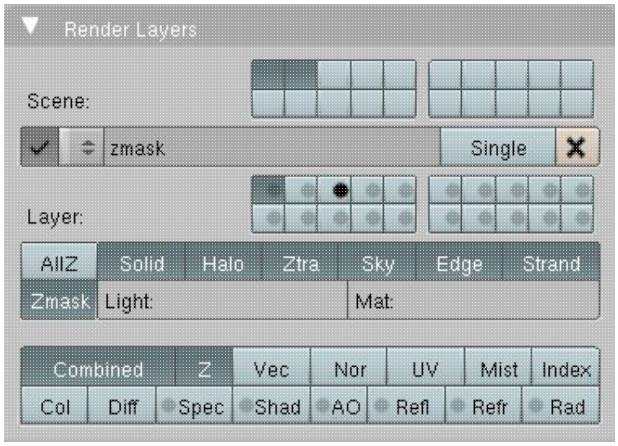
Setting up the correct field order
Blender's default setting is to produce Even fields *before* Odd fields, this complies with European PAL standards. Odd fields are scanned first on NTSC.
Of course, if you make the wrong selection things are even worse than if no Field rendering at all was used!
If you are really confused, a simple trick to determine the correct field order is to render a short test animation of a white square moving from left to right on a black background. Prepare one version with odd field order and another with even field order, and look at them on a television screen. The one with the right field order will look smooth and the other one horrible. Doing this simple test will save you *hours* of wasted rendering timeâ←¦

Fields and Composite Nodes
Nodes are currently not field-aware. This is partly due to the fact that in fields, too much information is missing to do good neighborhood operations (blur, vector blur etc.). The solution is to render your animation at double frame rate without fields and do the interlacing of the footage afterwards.

Renderlayers are used to separate your composite image into layers. Use RenderLayers for a specific reason - such as creating depth of field, relighting isolated elements within the image via a normal pass, adding a colorcast to specific portions of the image, etc. The keyword here is isolation. Renderlayers allow you to dissect, effect and or correct individual elements or groups within your composition before outputting your final render. This saves you from endlessly re-rendering your scene just to find out whether a correction is going to work or not.

This page is being updated IN ANTICIPATION of some of the changes expected in the NEXT release of Blender.
Specific featurees and options discussed may not be available in the current release.

Finding Render Layers



In the Buttons Window, specifically the Scene Render buttons, the first panel has two tabs: Output and RenderLayers. You're close, if not already there. Click the RenderLayers tab header. This is where you select the layers that you want to render, and the settings for the upcoming render.

By default, the active layers in your scene are also active here. You know this because you have acutely observed that there are TWO sets of layer buttons here: one above the Render Layer name, and another set below (that has all the buttons selected). However, you can change that set of layers here; simply (â § ShiftLMB to select multiple layers that you want to render, or LMB to select a single layer to render.

Layer Sets AND each other
Only the objects in layers that are selected BOTH in the main Scene Layer group AND the Render Layer Layer group will be rendered. So, if the Scene has only Layer 1 selected, and your Render Layer set specifies to render only Layers 2 and 3, nothing but the Sky (if selected) will be rendered.

You might also have noticed that there are *three* possible states for these buttons. In the picture, the first button is dark gray (selected), the second is light gray (not selected), but the *third* button has a black dot. This is used to indicate a **Z-mask**, as we will describe below.

Watch those black dots!
As you'll see, the idea of "black-dotted buttons" is used in two places on this panel, for two entirely different meanings. In *this* group of buttons, the black dots indicate a Z-mask. But as we will see on the next page, "black dots" are *also* used to indicate render outputs that should be excluded from the "Combined" output. Same visual effect; two meanings.

Using Render Layers

What are Render Layers *really* used for? Blender's node-based compositing system!

In Nodes, when you added an Input Node of type *RenderLayers,* and selected the Scene, you bring in whatever information you've specified for that RenderLayer. This node becomes a source for the rendering pipeline products you've specified *(see below),* as applied to the objects in the qualifying layer(s). Each of these products then "flow out of" that node toward their appointed destinations in the node graph you've constructed.

Render Layer settings

Render Layers are set up via four main groups of options:

- Enabling and Naming
- Scene Layers to Render
- Pipeline Products to deliver
- Passes to make available

Enabling and Naming

The checkbox enables or disables the computation of the whole render layer. Enable only those layers you are working with to save time. The selector allows you to scroll through and examine existing render layers, or to add a new one. You may name your render layers by â § ShiftLMB into the name field. The "Single" button makes Blender render ONLY that render layer and none other. Enable this if you are working on a single layer to avoid rendering unwanted layers. The X deletes the layer.

Scene Layers to Render

The next group of controls are the layers to render. As noted above, this set of layer buttons are ANDed with the currently selected scene layers (shown at the top of the panel). If both corresponding buttons are enabled, the layer is fed into the render pipeline.

You can select that layer by LMB clicking the button. To select multiple layers, â § ShiftLMB click. (The dot in the button in this case turns *dark gray*.)

In Addition... You can select that layer to feed the **Z mask** (explained below), via CtrlLMB. (To indicate this, the button's dot turns *dark black*.) Deselect by CtrlLMB again.

Pipeline Products to deliver

Each renderlayer has its own set of major products to include in the rendering pipeline. To save time and give you control when working with passes, this set of buttons allow you to select which major products to render:

AllZ
: Z-values are computed for everything in view, not just those things that are rendered. When disabled, objects not included in the render have no ("infinite") z value

Solid
: Solid faces are rendered. All normal meshes are solid faced.

Halo
: Halo materials are rendered

Z-transparency
: Transparency may be Z-based or Ray-traced. If Z-based, enabling *Ztra* renders transparent areas with the z-value of what is behind the transparent area.

Sky
: Turning on Sky renders the sky, as defined in your material world settings. Otherwise, a black alpha transparent background is rendered

Edge
: If Edge is enable in the Output panel, objects in this renderlayer are given an outline edge. Turning on

Edge pulls in the Edge settings from the Output tab, and adds an outline to the objects. Edges have to also be enabled on the Output tab.

Strand

Strands are strings of static particles that are colored as part of the material settings; they look like strands of hair, or grass.

Z-mask

The image rendered is from the objects that are between the selected layer(s) and the Z-mask layers. In the example, the cube is on layers 2 and 3, and the grass in on layer 1. In the render layer which we have arbitrarily chosen to call "zmask", as shown in the picture above, layer 1 is selected and layer 3 is designated as the Z-mask (as indicated by the black dot). Therefore, only that part of Layer 1 which is in front of the object on layer 3 (the cube) is rendered.

Light

Enter the name of a light group, and the scene will be lit with only those lights. Usually, you use this to speed up draft renders, of a scene that has complicated lighting, by entering the name of a small group of key lights.

Mat

Overrides all material settings to use the name of the Material entered. Use this to speed up draft renders. Use the default material to check basic lighting.

Passes to make available

Render Passes (Combined, Z, Vec, etc.) are discussed on the next page.

Creating a new Render Layer

By default, there is 1-Renderlayer created for you, and it includes all layers, whether they are used in your scene or not. To add yet another Render Layer, click the yellow up-down selector and select Add New render layer. You now have two Render Layers to choose from, and the active one is shown in the window. Each Render Layer will have its own set of layers that are rendered (sort of makes sense now, doesn't it?).

For example, you might have a robot in a scene with a ground object, buildings, etc. If the robot is on visible-layer 5, you would create one render layer named "Robot" with layer 5 selected in both the Scene: and Layer: buttons.

You would create another render layer (maybe named "stuff") that had all other layers EXCEPT layer 5 selected in both the Scene: and Layer: buttons. Then, back in the Node Editor, you would create TWO input nodes of type Render Layer: one for the Robot Render Layer, and another for the other Stuff. Run both through a mixer and out to the Composite viewer to get the big picture.

Examples

Rendering only certain objects

For example, suppose you have added a cool halo to your robot and you want to quickly see what it looks like. Suppose your scene has boxes on layer 1, laser rifles on layer 2, the robot on layer 5, and lights and camera on layer 20, and they are all selected and visible in 3d view. If you want to render just your robot, and he is on layer 5, you click on the render layer 5 button (which is below the Render Layer name), de-select sky (so that the sky/horizon is not rendered) and select Halo. Presto! When you render, only the robot is rendered (quickly) and not all the other elements of your scene (like the boxes he is running in front of).

Outlining only selected objects

To render an image where only one or two of the objects are outlined, move those objects onto layer(s) separate from everything else. Create a 1-Render Layer for those layer(s) by selecting only those layers in the Render Layer layer set. Create another 2-Render Layer for the other stuff. Enable the Edge option for 1-Render Layer (remember to also enable Edge on the Output tab) and make sure it is de-selected (off) for 2-Render Layer. In the Node Editor, create two input nodes, one for each Render Layer. Mix the two images. Done. Simple. Yea.

Render Passes are the different things the Blender Render Engine must calculate to give you the final image. In each 'pass' the engine calculates different interactions between objects

RenderPass In Detail

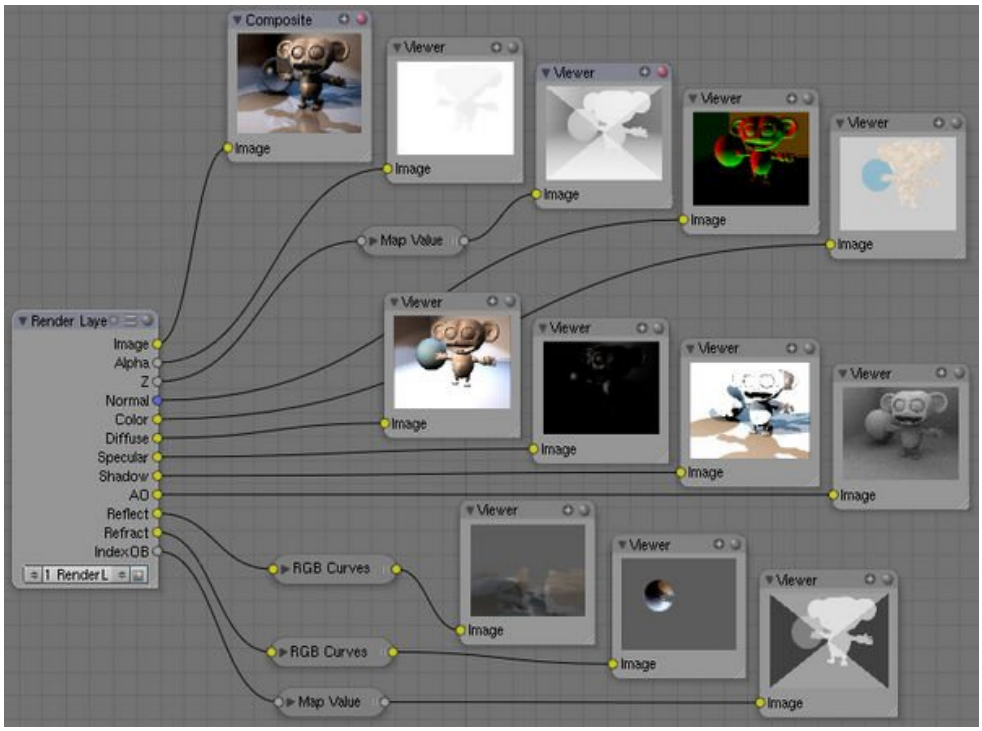
Everything you see in a render must be calculated for the final image. All interactions between objects in your scene, lighting, cameras, background images, world settings, etc. must all be separately calculated in different passes for different reasons, such as calculating shadows or radiosity - in a render, every pixel has been calculated several times to make sure it will show the right colour for the right part of the image. Various things that are calculated in a standard render include:

- *Where are **shadows** cast?*
- *How is **ambient** light in the environment blocked (**occluded**) by objects in the scene?*
- *How is light **reflected** off mirrored surfaces?* Like shadows, lines are calculated, except this time they come from the camera and bounce of mirrored surfaces, so that when these lines hit an object, the engine calculates that this is what the camera should see
- *How is light bent (**refracted**) as is passes through transparent objects?* Does it go straight through? Does it bend? If so, at what depth in the object?
- *What designated **objects** are in the scene, and what is their outline?* Should the object appear blurred, or should it appear in sharp focus?
- *How fast is something moving (**velocity**)?* Should it appear blurred with our framerate or is it slow enough to still be focused on properly?
- *How far away from the camera are objects' surfaces (**Z-depth**)?* Can the object's surfaces be seen at all, or are they being blocked by another object's geometry?
- *Does an object have a **normal** vector (bumpmap)?* Do shadows and apparent geometry need to be calculated for any objects?
- *Is there any **specularity**?* Are objects with textures such as metal shiny at all?

Renderer Rewrite

Starting with Blender 2.42, the render engine was rewritten. See Unified Renderer if you are using an old version of Blender.

The answer to each of the above questions is an image or map, shown below:

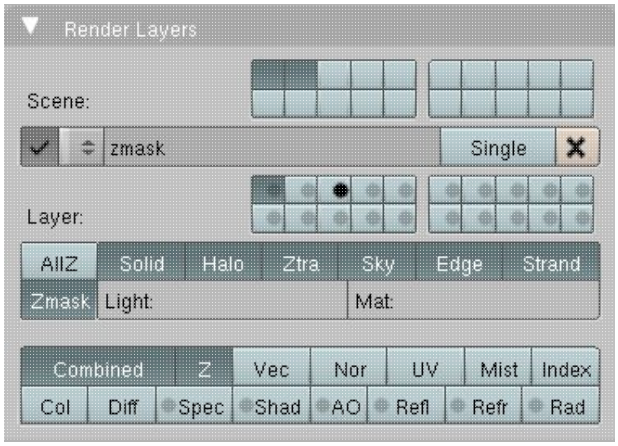


Each RenderPass puts out an image or a map. For the purposes of this example, a RenderLayer was defined to produce all possible outputs. When a RenderLayer input-node was added to the node diagram and the RenderLayer input-node was subsequently associated with the RenderLayer, all of the layer's outputs appeared as connection points on the right-hand (output) side of the node.

RenderPasses that produce Images can be directly viewed in a viewer, or, if they are the only pass that is rendered, saved as the render image. If the pass is enabled, it can be saved in a multilayer OpenEXR format.

If the RenderPass output is not an image but is a map, it needs to be translated into something that we can see. For example, the Z-depth map is an array of values that specifies how far away from the camera each pixel is; values range between +/-3,000,000 blender units or so. The intermediate node you see above, between the RenderLayer output socket and the Viewer node input socket (such as Map Value) does this translation or scaling. You must use that specific kind of translation node to get good results if you intend on operating on that map as an image. You must then, after making any adjustments, run the map back through that node to re-scale it back to original before saving.

Selecting RenderPasses



RenderPasses are the various distinct outputs that the renderer is able to generate. All of the following render outputs are normally combined into a single output known, appropriately enough, as the **Combined** output. But you can also select any of them to be output as a separate pass. (If you do so, in most cases you can choose whether to *also* continue to include it in the Combined output.)

Some of these outputs must be enabled and used within your scene (and not just selected in the RenderLayer panel) in order to show anything. For example, if you do not have any lights in your scene, or those lights have been set to not cast shadows, or objects in the limelight do not have materials which have been set to receive shadows, the **Shadow** pass will be blank; there's simply nothing to show you. If you have not enabled *Ambient Occlusion* in your World environment settings, the **AO** pass will be blank, even if you select it here.

To save time and disk space, you have to tell Blender each of the passes to render in the RenderLayers panel (which we first introduced on the previous page):

- **Combined**: This renders everything in the image, even if it's not necessary. ("The whole enchilada" so to speak.) This is all the options below, blended into a single output, *except* those opions which you've indicated should be omitted from this pass, as indicated by a "black dot" within the button. (See below.)
- **Z**: the Z-depth map; how far away each pixel is from the camera. Use for Depth-Of-Field (DOF). The depth map is inverse linear *(1/distance)* from the camera clip start.
- **Vec**: Vector; how fast things are moving. Use with Vector Blur
- **Nor**: Calculates lighting and apparent geometry for a bumpmap (an image which is used to fake detail on an object) or for changing the apparent direction of light falling on an object.
- **UV**: allows texturing after rendering. See UV node.
- **IndexOb**: masks selected objects. See MaskObj node.
- **Col**: the basic colors adjustment in the image
- **Diff**: Diffuse; main colors before being shaded
- **Spec**: Specularity adjustment
- **Shad**: Shadows cast. Make sure shadows are cast by your lights (positive or negative), and received by materials. To use this pass, mix multiply it with the Diffuse pass.
- **AO**: Ambient Occlusion. Make sure it's turned on in your environment and that RayTracing is enabled.
- **Refl**: Reflection off mirrors and other reflective surfaces (highly waxed white floors, for example). Mix Add this pass to Diffuse to use it.
- **Refr**: Refraction of colors through transparent meshes. Mix Add this pass to the Diffuse pass to use it.
- **Rad**: Radiosity; colors emitted and cast by bright objects spill onto other objects

When you enable a pass, the appropriate socket on the RenderLayers node shows up like magic, and can be used as shown in the example above.

As we said, the **Combined** output is an amalgam of several outputs which are *also* available separately. When you select one of these outputs, they will be provided separately *and also* included in the Combined pass. You

can omit outputs from being included in the Combined pass by Ctrl+Click on the appropriate buttons, which will cause a "black dot" to appear in the button. Where the black button appears, outputs will *only* be provided separately: they will not also be included in the Combined pass.

Watch those "black dots!"
The visual effect of "buttons with black dots" is used *twice* in this panel, with two very different, *unrelated*, meanings. (On the previous page, they refer to Z-masks.)

De-Selecting Render Passes

When you CtrlLMB CTRL-click on a pass, the particular pass will be excluded from the combined pass. Excluded passes are marked by a black dot next to their name. They will be made available separately *but not* included in the combined pass.

(After all, sometimes when you say you want "the whole enchilada," you really don't want the *whole* enchilada!)

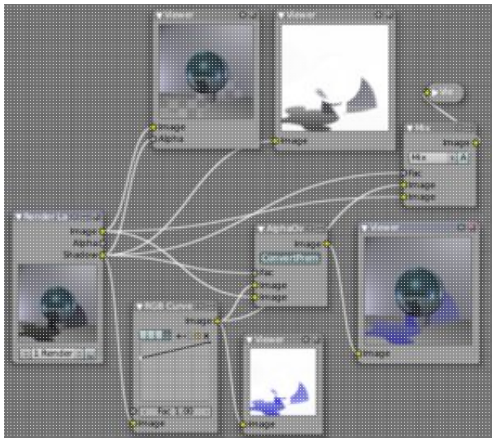
Using RenderPasses

The primary purpose of RenderPasses is to enable you to process the various outputs in different ways, by constructing networks of render nodes. You can achieve many special effects, and economize considerably on the render times of complicated scenes, by creative and effective use of this facility. We'll show you a few examples of this in just a moment.

Quite a bit of information about the typical uses for some of the passes is discussed elsewhere:

- Image: Since this is the main product, all of Blender uses it.
- Alpha: See the *AlphaOver* node and all of the *Matte* nodes
- Z: See the *Defocus* node
- Vec: See the *Vector Blur* node
- Normal: See the *Normal* node

Recoloring Shadows



Let's run the Shadow buffer through a colorization noodle, then recombine it and all your shadows will be artificially colored. Lots of threads in this noodle shown to the right, so let's walk through it. On the left is the RenderLayer input node: it refers to one of the RenderLayers that we have defined for our scene. In the scene, we have a reflective ball on a pedestal standing in front of a backdrop. Everything (except the ball) is gray. We use a standard four light rig: backfill placed high, two sidefills at ground level, and a key light above and to the left of camera. Suzanne, a monkey shaped geometry, is standing in front of the key light, so her shadow is cast into the scene on the floor. The ball casts shadows onto the backdrop and floor.

The output channels of the RenderLayer node are determined by which buttons we selected when defining our RenderLayer. The top two viewers show you the image output using the Shadow as the Alpha channel, and the node next to it just the Shadow channel. Where the Shadow is dark, the image in the left viewer is transparent. We have used the Shadow to cut out parts of the image.

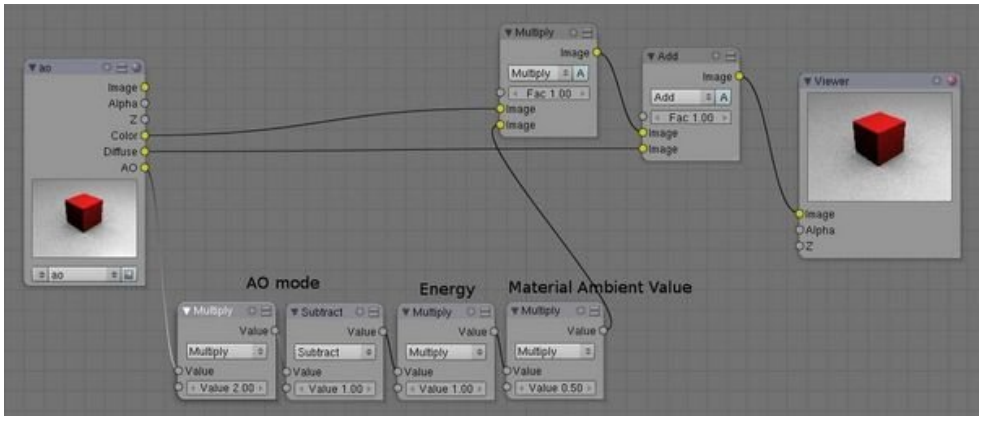
We then take the shadow through an RGB Curve, which is set to magnify just the Blue by 75%; so a gray shadow of (R:40, G:40, B:40) becomes (R:40, G:40, B:40x1.75=70). That blue-tinged shadow is shown in the bottom viewer. Now we have two options: AlphaOver and Mix. For either option:

- Use the Shadow map as a Factor
- Feed the Blue Shadow to the Top Socket
- Feed the core or base image to the Bottom Socket

The resulting image is the same in either case; a blue shadow. Note that Suzanne's reflection is not blue; there's a different RenderPass for that.

You could just as easily swap in another image entirely; for example, the shadow map from another renderlayer. You can even take an image from another project entirely and use that instead (using the Image Input node), to get a different effect. (For example, an effect similar to a *Star Wars Episode One* movie poster, where Anakin Skywalker already casts the shadow of Darth Vader.)

Compositing Ambient Occlusion



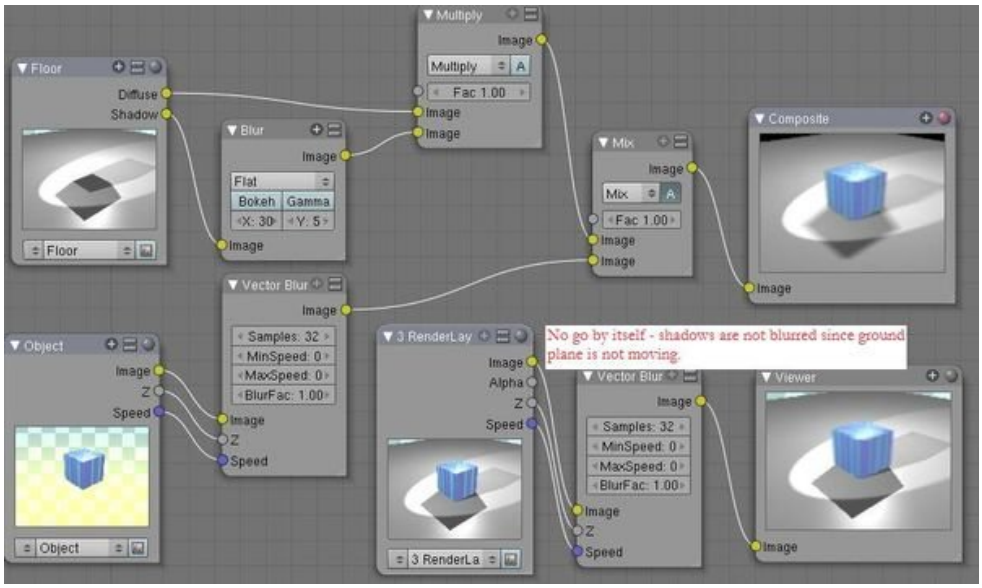
AO is a geometry-based dirt shader, making corners darker. It is separately enabled in the World settings and computed as a separate pass. When enabled, it has one of three Modes (*Add*, *Subtract*, *Both*), and variable *Energy* level (which changes the intensity of the shading). The third variable is the amount of Ambient light that the material receives. If it does not receive any, then ambient occlusion does not affect it. Based on these variables, Blender computes an AO pass. If you call it out as a separate pass and wish to composite it back into your image, you will need to enable the Color and Diffuse pass as well.

To configure your noodle, consider the example image above.

1. . First, depending on the AO mode do one of the following: If AO mode is Add: directly use the AO pass. If AO mode is Sub: Calculate AO - 1, or if AO mode is Both: Calculate 2*AO - 1
2. . Multiply the output of Step 1 with the AO energy level
3. . Multiply the output of Step 2 with the material's ambience value. If you have materials which receive different ambience light levels (0.5 is the default), one would have to create an ambience map based on Object ID)
4. . Multiply the output of Step 3 with the color pass
5. . Add the output of Step 4 to the diffuse pass

If shadows, colored ambient light, specularity, reflections, and/or refractions are involved they have to be added to the diffuse pass before adding the converted AO pass.

Vector Blurring Shadows



When using Vector Blur, instead of Motion Blur, objects in motion are blurred, but objects at rest (with respect to the camera) are not blurred. The crossover is the shadow of the object in motion. Above, we have a cube in motion across a ground plane. If we just ran the combined pass through Vector Blur, you can see the result in the lower right-hand corner; the box is blurred, but its shadow is sharply in focus, and thus the image does not look realistic.

Therefore, we need to separate out the diffuse and shadow passes from the floor by creating a "Floor" renderlayer. That renderlayer has Diffuse and Shadow passes enabled, and only renders the floor object (layer 2). Another renderlayer ("Cube") renders the Z and Vector passes, and only renders the cube (on layer 1). Using the Blur node, we blur the shadow pass, and then combine the diffuse and blurred shadow by multiplying them together in a Mix Multiply node; we then have a blurred shadow on a crisp ground plane. We can then mix the vector-blurred object to provide a realistic-looking image.

Conclusion

Render Passes can be manipulated from Blender 2.43 to give you almost complete control over your final image. Causing objects to cast shadows that aren't really their shadows, making objects appear out of focus or sharply in focus like a real camera, manipulating colours just for final post-processing or just reconfiguring your render passes to save render time, are all things which you might wish to manipulate the render engine for.

Happy Blending

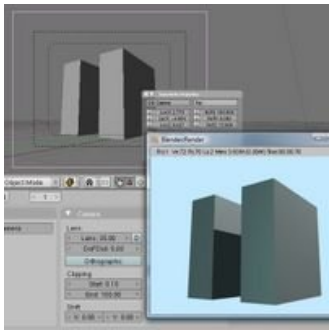
Vanishing points in Render

When you press F12 and get your render, you see an image as seen through the camera's "perspective". Like how you can *view* your model in 3D View from the top, front, side, or user perspective, you can *render* your object from different perspectives. This perspective takes into account the lens size, type, and offset in giving you that picture. Each perspective uses a different number of vanishing points. If you look at a 3D image of a cube, you will see three kinds of edges: vertical, horizontal, and depth. If all of the vertical edges are exactly parallel, there is no vanishing point for them. If however, they are not parallel, if you extended them by continuing them with a ruler, they would at some point intersect. That point is called the vanishing point.

For special purposes, different kinds of render cameras can be set up to give you different perspectives. For reasons discussed below, you may wish to limit the number of vanishing points, especially for architectural purposes. Architects and drafting people are responsible for rendering the object or building with true dimensions and true relative proportions. If you look at that example render, the building looks all sorts of distorted, like it had been made of mud and was collapsing. If you told a builder to build that, you would end up with a building that actually had leaning walls and rooms that were narrower at the top. Way back in the old Greek days, when they started building tall columns, they built them thicker at the top than at the bottom, so that when viewed looking up, the two sides would look straight up and down. Then they even started narrowing the columns at the top to give the illusion that the building was taller and would look higher. During the Renaissance, the concept of using vanishing points in art evolved. Blender offers a few tricks of its own to let you do the same.

Note
To follow sections below you will need to know how to adjust Camera Settings

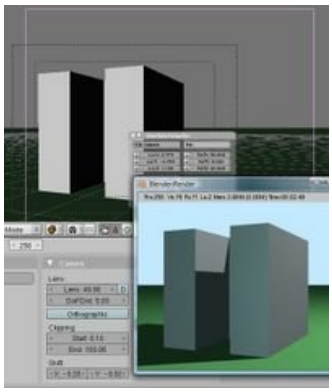
Three Point Rendering



Normal Three Point Render
When looking at or rendering a picture of a high building from ground level off to one side, and aiming up, using the normal 35mm camera, you get 3 point perspective. If you laid a ruler along the vertical lines, you would see that they converge to a point above the building. The horizontal lines are converging off to one side (the left in this example), and depth (receding) lines are converging to a different third point (somewhere off to the lower right in this example). Hence the name 3-point rendering - there are three vanishing points.

This is reality, and there is nothing wrong with that. When you next step outside and look at a tall building, this is what you actually see. However, your mind knows that the building is square, and can adjust your perception of the building so that you are not scared that the building is going to fall over.

Two Point Rendering



Two Point Vertical Render

Normal architectural rendering is called two point rendering; when vertical lines are parallel, and horizontals, if followed out to the side, converge on one point, and receding or depth lines converge to a second point. Architects often like this Two-Point rendering, so that the sides of their buildings are completely vertical and don't appear to be falling inward. This is also quite nice for compositions and schematics, given that the lines of the paper you print on and the screen you view with are also straight. Previously to get a 2-point perspective, you had to aim the camera level to the horizon, however this resulted in the top half of the building being cut off and the horizon being in the exact middle, which looks very boring. Architectural photographers use 'shift lenses' to solve this problem. Shift lenses shift the image to another place on the film.



Two Point Horizontal Render

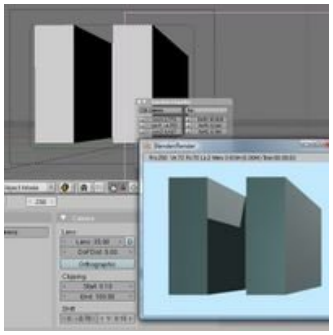
This technique works well for high buildings as well as for normal sized objects. Most of the time, the two vanishing points are horizontal and depth lines, with the vertical lines parallel. However, some titles are done with the horizontal lines parallel, and the vertical and depth lines having the vanishing point. This dramatizes and exaggerates the massiveness and height of the title. To get this effect, position the camera at ground level, centered, angle the camera upward, and shift the render passpartout down. In the example, the camera is rotated 30 degrees upward, at ground level with the title. A bright key light with a short falloff provides dramatic lighting that is bright in the middle and falls off toward the sides, further enhancing the depth.

To achieve 2-point rendering:

- Use a short wide angle lens camera, say with a Lens Size of 10 mm placed close to the building, or a long lens farther away from the building. These differences affect the depth of the building render, with longer lenses making the building appear thinner and less dramatic or distorted. The example uses a 40mm lens.
- Position the camera off to one side of the object, vertically halfway up the building to minimize distortion of the vertical building edges. You may alter this vertical (Z value) position to be slightly higher than ground level or higher than the top (if you want to see the top of the object or building). To show the front bottom corner of the building jutting out, raise up the camera.
- Angle the camera to be looking away from the building and directly level at the horizon - not pointed up or down (note the 20 degree Z angle in the example). This should make the vertical lines parallel. The more the camera looks at the object, the closer the vanishing point for the horizontal lines, and perceived depth will increase as that vanishing point gets closer as well.

- You may have to angle the camera slightly down (just 1 degree or so) so that vertical lines appear vertically up and down, both near and far. If the lines are curved, use a longer lens. With your 3D View set to Camera view, use the passpartout or pixels on your monitor to determine vertical.
- Move the camera toward/away from the object until it appears near a corner of the render and is the right size.
- Adjust the Shift: X and Y settings until your object is positioned properly.

One Point Rendering



One Point Render

One point rendering is where vertical and horizonal lines are parallel, and depth lines converge at one point. Architects really like these renders, since the front-facing faces are true and square, and the building recedes off into the distance so that it looks like it has some depth. If the camera is placed at ground level, even with the bottom of the building, it really looks dramatic but orderly in a weird sort of way. Title graphics are sometimes rendered this way.

To get 1-point (1pt) renders,

- To get more dramatic depth lines, use a short wide angle lens camera, say with a Lens Size of 10 mm, very close to the building. For a more normal appearance, stick with the 35mm lens.
- Position the camera off to one side of the object, slightly higher than the top (if you want to see the top of the object) or at ground level (the example image has the camera almost at ground level). If you position the camera *below* ground level, the bottom depth lines and horizontal lines will merge up (become congruent) for a *very dramatic* effect.
- Angle the camera looking straight back, perpendicular to the true face. Vertical lines should be parallel. Rotate the camera on the Z axis *slightly* toward the object until the horizontal edges are also parallel. Technically, you are correcting for parallax (just a casual line to drop on your girlfriend to impress her). The example has the camera rotated 0.5 degrees toward the object.
- Move the camera toward/away from the object until it appears at the proper size relative to your passpartout.
- Adjust the Shift: Y settings until the bottom of the passpartout (or title line if you want to show some approach ground in front of the building) is even with the bottom of the building. Adjust the X setting until the building is centered (or slightly offset from center for artistic appeal, or to show the parking lot next to it) as shown.

In the example screenshot, the Lens is 35, X is negative and Y is positive. The camera is off to the right of the object, even with the bottom of the building. If X & Y were zero, the building would have appeared off camera, in the upper left-hand corner of the passpartout.

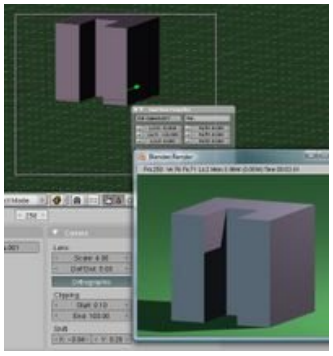


Parallel Horizontal Edges:

You can use the lines of the passpartout as a guide in rotating the camera to determine when the

horizontal edges are parallel.

Zero Point (Orthographic) Rendering



Orthographic Render

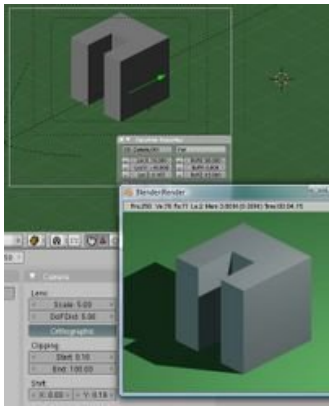
Zero point rendering is where vertical, horizontal AND depth lines are all parallel, and is commonly rendered at 45 degree, 30 degree, or 60 degree angles. With all of those sets of edges parallel to each other within that set, there are no vanishing points. The example shows that same building rendered at 45 degrees from all angles. Note that the vertical lines are parallel to each other, the horizontals, and the depth lines are parallel to each other. From this, it is very easy to see that the left top edge of the building is the same length as the right top edge, and that the building is as deep as it is wide and high; if you measured the edges with a ruler, they would all be the same. Orthographic rendering gives a true mathematical render of the shape of the object. An Orthographic perspective is what you see in the User View of a 3D window (if View->Orthographic is turned on).

To get an Orthographic render:

- Enable Orthographic in the Camera panel. This makes at least one face to be true to the camera.
- Point the camera at the object
- Position the camera or alter the Scale so the object is the desired size

With Orthographic cameras though, Lens size is irrelevant, since light rays do not converge to the camera from a field of view. They come in parallel, and so you can only Scale the camera size to take in more or less of that huge plane. Note that Shift X & Y are zero, and that the camera is positioned perfectly off at a 45 degree angle to the object/building, and is rotated exactly 45 degrees to face the building. Thus, the near edge is aligned with the back edge (since the object is square). Orthographic renders are usually made at 30, 45, or 60 degree angles to the object. Specific measurements are left to reader using triangle math.

Isometric Rendering



Isographic Render

While we are at it, we might as well cover Isometric rendering, which is a very specific type of orthographic render very often used in drafting and third-person computer games. In Isometric renders, you want your depth lines and your horizontal lines to be at 30 degrees off horizontal, and your vertical lines to be, well, vertical. Some complicated vector calculus in the wikipedia gives us a convenient shortcut. To get Isometric Renders:

- Make your camera Orthographic
- Add a "Track To" constraint (Object F7 context, Constraints panel) to the camera for it to Track To the object (type the name in the Target OB: field), using To: -Z and Up Y.
- Position your camera so that it is 45 degrees in the XY plane from your object, and raised at a 30 degree angle. If your object is at XYZ (0,0,0), then your camera should be at (10, -10, 10), or for a view from the left side, (-10, -10, 10)
- Adjust the Scale of the camera (Editing F9 context, Camera panel) so that the object fits within the passpartout
- Adjust the Shift: Y value so that the object is centered in the render.

Depth Of Field (DOF) Explained

Real world camera lenses and your eyeball transmit light through a lens (cornea) that bends the light, and an iris that limits the amount of light, to focus the image onto the film, CCD/Cmos sensor, or retina. Because of the interaction of the lens and iris, objects that are a certain distance away are in focus; objects in the foreground and background are out of focus. We call this distance their depth, or â← Zâ← distance from the camera or eye.

Light comes to the lens (in the real world) at an angle; from some direction. What you see depends on your perspective; if you move closer, different angles of the scene are revealed. To make â← flatâ← pictures, like an architectural drawing or plot, Blender can also make an orthographic rendering. So, there are two kinds of renderings, Perspective and Orthographic. Perspective simulates light coming in at an angle to the lens from the field of view, and Orthographic (disabled by default) simulates light coming straight in to an infinitely large backplane or flat retina.



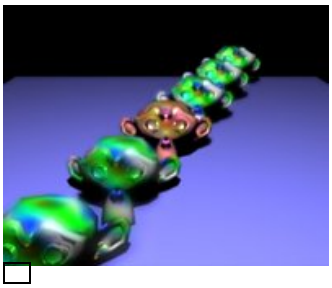
Depending on the diameter of the iris, there is a range (of distance) where objects are in focus. In cameras, the diameter of the iris is controlled by an â← f-stopâ← . Said another way, there is *field* of view that you see left to right, up and down; your â← pictureâ← , if you will. At a certain range, or *depth* away from your eye, things are in focus. For example, at night, you may be able to focus your eye on objects that are 10 to 15 feet (3 to 5 meters) away. Anything closer than 10 or farther away than 15 is blurry. Your **depth of field** is thus 5 feet (2 m).

The larger the iris, the smaller the depth of field. This is why, during the day, you can focus on a range of things stretching out far from you. In film, there is a person whose job is to measure the distance from the camera to the actorâ← s nose, to ensure that the focus is set perfectly.

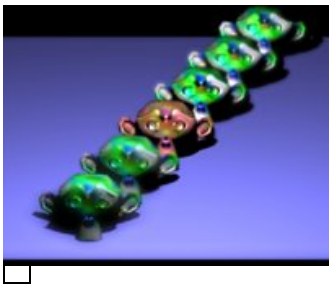
The more that an object is out of its depth (the perfect value for this depth is called *focal plane*), the blurrier it is. In fact, the depth of field is the range on both sides of the focal plane in which the blurriness of the objects is considered to be low enough to be imperceptible. In Blender, this distance is called the Dof Dist or â← Depth of Field Distanceâ← and is set in the Editing context (F9) for the camera. Alternatively, you can have the camera automatically stay focused on an object, by entering the name of the object in the Dof Ob field.

Field of View and Lens Size

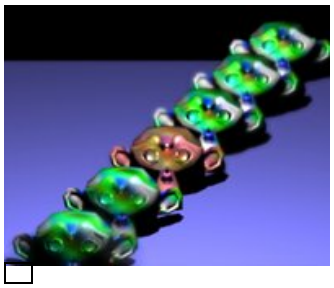
The field of view varies by the size of the lens. With cameras, a 35mm lens is kind of a standard size because the picture it takes mimics the size of the picture seen by the eye and pictures can be taken rather close. In Blender, use the Camera settings to change the size of the lens (35mm is the default). A longer lens taking a picture farther away has the same field of view, but has a different perspective of the view that many directors love because it â← condensesâ← the scene and smooths a sweep, since it is farther away from the action:



35mm lens from 10 units away.



210mm lens from 60 units away at same location/rotation.



210mm at 50 units; repositioned to frame the view similar to the 35mm shot.

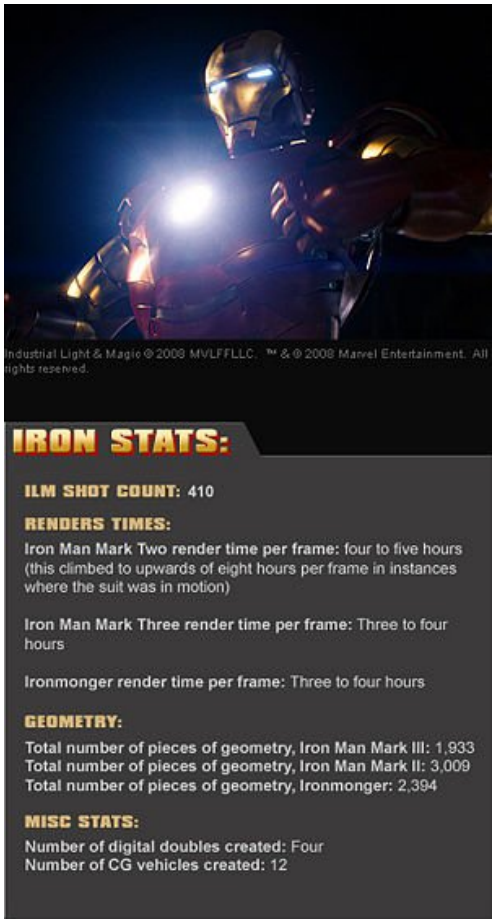
Zooming in Blender

Zoom is the ability to expand a subset of the picture; we humans have no such ability. Well, I take that back; we do: we just get up off the couch and walk up closer to what we want to see (however, this is more like â← travelingâ← than â← zoomingâ←). Blender allows you both actions: you can move the camera closer to or farther away from an object for a track (or â← truckâ←) in/out, and/or change its lens size. You can automate these by assigning an Interpolated (Ipo) curve to the object or to the camera, respectively.

Depth of Field in Computer Graphics

In computer graphics (CG), there is no physical lens or iris, so the depth-of-field (DOF) is infinite and all objects are always in focus. However, for artistic reasons, we want our main characters to be in focus, and everything else a little blurry, so that our audience does not focus on distracting things in the background. Also, it is easier to discern the main actors when they are in focus, and everything else isnâ← t. So, we have to create an effect, or **Depth of Field Effect**, to composite our images and post-process them to achieve realistic-looking results.

Rendering Big Hairy Scenes or Speeding Up Rendering



"A watched pot never boils" is the old saying, but you may wonder why your render takes so long to create, or worse, crashes mid-way through! Well, there is lots going on and lots you can do to speed up rendering or enable a complicated render to complete. Also, it is possible to render a very complicated scene on a mediocre PC by being "render-smart". Here's a "top ten" list of things to do or not do in order to speed up rendering or even avoid crashes on scene render. Some options may decrease the quality of your render, but for draft renders you may not care.

If you get the message "Malloc returns nil", in plain english that means the memory allocator tried to get more physical memory for Blender but came back empty-handed. This means that you do not have enough memory available to render the scene, and Blender cannot continue. You will need to do one or more of the following tasks on this page in order to render.

Hardware Improvements

1. Get more RAM up to your PC's (motherboard and operating system) limit. Presently, Blender can use up to 8GG (giga-giga) of physical memory (64-bit address space), but most PC's can only handle 2G of RAM.
2. Upgrade your CPU to a multi-core/multiprocessor
3. Upgrade your OpenGL video drivers
4. Get a faster memory bus
5. Get faster memory, up to yor PC's motherboard limit. 667MHz memory is 30% slower than 800MHz.
6. Use or set up a render farm using all available PC's in your house, or use a render farm such as BURP.

Operating System Configuration

1. Increase the processing priority through your OS
2. Increase your swap file space used by the OS for memory swapping. Also called virtual memory pagefile size, up to the size of your physical memory.
3. Upgrade to a 64-bit operating system
4. Exit or stop any background processes, like virus scanners, BOINC, Real, even "inactive" ones like Quicktime, as they can randomly start up to look for updates. They also take up RAM.
5. Disable network connections to stop random pinging traffic and refresh traffic
6. Stop listening to internet radio and close web browsers, especially any multi-media (audio/video/game) sites.
7. Close down all other running applications, like Word
8. Exit all TSRs (those icons on your system tray) and widgets, any background processes, and virus scanners

Blender Choices

1. Increase the MEM Cache Limit in the User Preferences System & OpenGL tab
2. Upgrade to an optimized Blender build, especially if you have a modern chip that supports SSE2 - render times are **30% faster** using an optimized build.
3. Switch to an Orthographic camera, and render your own "parts" of the scene as separate images, and then paste those parts together in GIMP. An old trick in making your own panoramic with a real camera is to take three or so pictures of a very wide (beach sunset) scene, where you take one picture, rotate to the right, snap another, then another, and when you get the pictures developed, you overlap them to make a very wide landscape image. Do the same in Blender; render out one shot to a file, then move the camera to look at a different area of the scene, and render that shot. Each shot will be of a smaller area and thus take in fewer polygons/faces. Be sure that when you position your camera that you snap overlapping shots, so that you can then match them up. If you don't want to use GIMP, you can use compositing nodes and the Translate node to match them up in Blender.
4. Minimize the render window (and Blender if rendering to an internal window). ATI users report dramatic speedup on a per frame basis, which adds up over frame range.
5. Use the Big Render script to render sub-sections of the overall image, and then paste them together.
6. Make a customized build. For example, comment out calls to check_non_flat_quads in convertblender.c, in some cases can make a noticeable difference 20-40%

Scene and Specific Objects

1. Remove lamps, or move them to un-rendered layers, or tie them to layers
2. Turn off some lamp's shadows, using only one or two main sun lamps to cast shadows. A few "shadows only" lights will render faster than every light having shadows on.
3. Use Buffer Shadows rather than Raytraced Shadows
4. Bake your shadows using Render Baking Full Render bake on surfaces that do not move. Using that texture for that mesh, then disable shadows for that material.
5. Simplify meshes (remove polygons). The more vertices you have in camera, the more time it takes to render.
6. Remove Doubles, or use the Decimator mesh edit feature.
7. Remove Subsurf and Multires modifiers
8. Delete backsides of meshes (removing unseen vertices)
9. Render just a few objects at a time; in the beginning of your project, render the background objects and set that will not change and will always be in the background.
10. Put the buildings on another layer, and through renderlayers, don't render them. Then composite them back in later.

11. Make the camera static so that you can better accomplish the above two ideas.
12. Avoid use of Area lights
13. Make materials Shadeless
14. Render Bake AO and textures, and then make those materials Shadeless
15. Decrease the Clip distance for spot lights
16. Decrease the Clip distance for the camera
17. Turn off world AO
18. Turn off Material SSS
19. Use smaller image textures. A 256x256 image takes only 1% of the memory that a 2k image does, often with no loss of quality in the ulitmate render.
20. Reduce Subsurf. Each level quadruples (4x) the number of faces from the previous level.
21. Reduce Multires
22. Make a matte render of background objects, like buildings, and put the image of them on a billboard in the scene instead of the object themselves. This will reduce vertex/face count
23. if you have lots of linked instances of an object - use DupliFaces, as these are instanced. If you have 100 of them, Blender will only store the geometry for 1 (Instances themselves take a small amount of memory).

Render Settings

- Output Panel
 1. Disable Edge rendering.
 2. Save Buffers
 3. Render to an Image Editor window, not a popup Render Window.
 4. Use multiple Threads on a multi-core CPU (with multiple Parts).
- Render Layers Panel
 1. Render only the Layers of interest.
 2. Render with all lights set to one simple spot (enter its name in the Light: field)
 3. Render with one material override (enter its name in the Mat: field)
 4. Disable unnecessary Render Passes, such as Z, or only render the pass of interest, such as Diffuse.
- Render Panel
 1. Turn off Shadows
 2. Turn off Environment Mapping
 3. Turn off Panoramic Rendering
 4. Turn off Raytracing
 5. Turn off Radiosity
 6. Turn off SSS Subsurface Scattering
 7. Turn off or lower oversampling/aliasing OSA
 8. Turn off or lower Motion Blur
 9. Render in Parts. This will also allow you to render HUGE images on a weak PC. On a multi-core PC, it will assign a thread to each part as well.
 10. Increase the octree resolution
 11. Render at a percentage size of your final resolution (like 25%)
 12. Turn off Fields rendering.
 13. Use Border rendering to render a subset of the full image.
- Anim Panel
 1. Decrease the frame count of the animation (and use a lower framerate for the same duration of animation). For example, render 30 frames at 10 frames per second for a 3-second animation, instead of 75 frames at 25 frames per second.
- Bake Panel
 1. Bake Full Render - create a UV Texture that colors the objects based on materials, and then use that UV Texture shadeless instead of the material - 7x faster:

2. Bake Ambient Occlusion only
3. Bake textures for objects
4. Baking Normals or Displacement does not speed up render time, and are used for other things.

- Format Panel
 1. Render at a lower resolution. Smaller pictures take less time to render.
 2. Choose a faster CODEC or CODEC settings
 3. Render in black and white (BW button)
 4. If using FFMPEG, do not activate Multiplex audio
 5. If using FFMPEG, Autosplit Output (Video panel button)
 6. Render only RGB if you need color; the A channel (RGBA button) takes more memory and is unused when saving a movie file.

Another Way To Do It: Compositing

Another strategy that can be used to address the problem of long (re-)render times is to structure your workflow from the ground up so that you make aggressive use of *compositing,* as described in the "Post-Production" section. Here, you seek to break down each shot into components that can be rendered separately, then combined later to create the finished clip. For instance:

- If the camera isn't moving, then neither is the background: only a single frame is needed. The same is true of any non-moving object within the frame.
- Both shadows and highlights can be captured separately from the objects that are being illuminated or shadowed, such that the intensity, color, and depth of the effect can be adjusted later without re-rendering.
- Tricky lighting situations can be avoided by handling the objects separately, then combining the individually-rendered clips and "tweaking" the result.

There are sometimes compelling advantages to this approach:

- In practice, the deadline killer is *re*-rendering, which ordinarily must be done (in its entirety) just because "'one little thing' about the shot is wrong." Compositing helps to avoid this, because (ideally...) only the specific parts that are found to be in error must be repeated.
- It may be that all you really need to do is to repeat the final compositing step(s). For example, to darken a shadow, change the color of a light, or adjust color timing, these changes can usually be done on-the-spot. By and large, these stages work *two*-dimensionally, manipulating what is by that time "a raster bitmap with R, G, B, Alpha *(transparency...)* and Z-Depth information," so they're consistently fast.
- Since each discrete rendering task has been simplified, the computer can carry them out using much less resources.
- The tasks can be distributed among several different computers ... even less-powerful ones (like the two older machines that are sitting in your closet right now because you can't get rid of them).

Of course, this approach is not without its own set of trade-offs. You must devise a workable asset-management system for keeping track of exactly what material you have, where it is, whether it is up-to-date and exactly how to re-create it. You must understand and use the "library linking" features of Blender to allow you to refer to objects, nodes, materials, textures and scenes in a carefully-organized collection of other files. You need to have a very clear notion, *in advance,* of exactly what the finished shot must consist of and what the task-breakdown must be. You must be a scrupulous note-taker and record-keeper. But sometimes this is the best way, if not the *only* way, to accomplish a substantial production.

Introduction

Compositing refers to building up your image or file from components. Like plywood, the final product is a combination of layers bonded together. The layers work together to create a better product than any single piece. Compositing has two main purposes: *Assembly* and *Enhancement*.

Image or sequence *assembly* is where the final footage is assembled from multiple shots of raw footage. For example, the two-second shot of the actor driving is taped to a three-second shot of the car driving by, resulting in a five-second sequence.

Image *enhancement* changes the appearance of a shot. In film, a shot may be taken during a bright and sunny day, but in post-production, the editor/producer may want the sky to be dark and cloudy to convey a sense of foreboding. The end result is thus a composite of what was originally shot, with effects layered on top and bonded to the original.

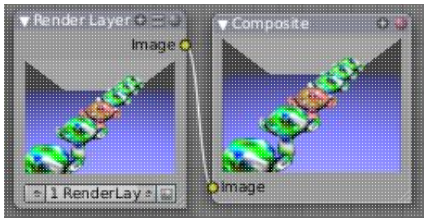
Compositing is accomplished two ways in Blender:

- Assembly
 - ♦ Using the Video Sequence Editor
 - ♦ Using Composition Nodes
- Enhancement
 - ♦ Using Composition Nodes

The Video Sequence Editor is used to *assemble* your shots into raw footage. The editor allows you to add a few mixing and transition effects, such as alpha over and wipe, and was the first way Blender enabled compositing. Limited post-processing can be done using the editor. This set of tools has been modularized and greatly expanded by the introduction of Composition Nodes. Composition Nodes are little workers that do something specific to a shot, such as colorizing it or sharpening it; you arrange these nodes and thread them together in assembly-line fashion to composite the images, resulting in processed footage.

To learn how to use the new, wild, and very cool Composition Nodes, click here. To use the Video Sequence Editor, click here.

Compositing Nodes allow you assemble and enhance an image (or movie) at the same time. Using composition nodes, you can glue two pieces of footage together and colorize the whole sequence all at once. You can enhance the colors of a single image or an entire movie clip in a static manner or in a dynamic way that changes over time (as the clip progresses). In this way, you use composition nodes to both assemble video clips together, and enhance them.



To process your image, you will use nodes to import the image into Blender, change it, merge it with other images, and finally save it. Think of nodes as workers on an assembly line. Each worker performs one basically simple task before passing the product on to the next worker. The assembly line is called a thread, and the image data enters and leaves a node by sockets. You connect one node to another by creating a thread that runs from one node's output socket to another node's input socket. Any number of threads can connect to the same input socket. A bunch of connected nodes is called a **noodle**. The example to the right shows the simplest noodle; an input node threads the camera view to an output node so it can be saved.

Compositing nodes are connected by threads into a map that routes an image through a sequence of nodes. Each node performs some operation of the image. You create the map using the Node Editor window type. Before you continue reading, it is elementary that you understand how to use this window type. Read all about the Node Editor by clicking here.

Term: Image
We use the term *Image* to refer to a single picture, a numbered sequence of images that, when played flipbook style, result in a movie clip, a single frame of a movie sequence, or a movie clip.

Every node has a similar set of controls and conventions. For more information on working with node controls, click here.

As with other cases where nodes are used within Blender, you can imagine that nodes fall into one of three general categories:

- **input** nodes, which introduce data into the noodle (*e.g.* from an image or layer). A noodle can have any number of inputs, and/or instances of the same input.
- **filter** (or **transformation**) nodes which process their inputs in some way to produce some outputs (that are intended to be used as inputs elsewhere).
- **output** nodes, which consume the data that is sent to them, perhaps displaying it or saving it to a file. A noodle can produce any number of outputs.

Any number of nodes can be connected in any number of ways, giving you tremendous control and creative opportunities.

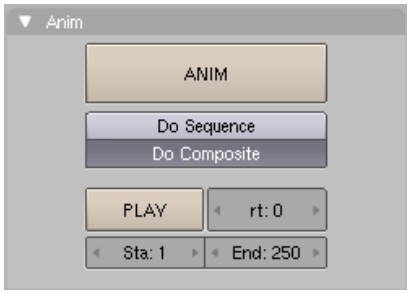
Accessing and Activating Nodes



Select the Node Editor window
Access the Node Editor and enable Composite Nodes by clicking on the correct icon (picturing a Face).



Node Editor Header with Composite Nodes enabled.
To activate nodes for the composition, click the Use Nodes button. Blender creates a default starting noodle, consisting of two nodes threaded together.



Use Composition Nodes
To use this mini-map, you must now tell Blender to use the Compositing Node map that has been created, and to composite the image using composition nodes. To do so, switch to the Scene button area (F10) and activate the **Do Composite** button located below the Animation button. This tells Blender to composite the final image by running it through the composition node map.

You now have your first noodle, a RenderLayer input node threaded to an Composite output node. From here, you add and connect nodes in a sort of map layout to your heart's content (or physical memory constraints, whichever comes first). There are many types of compositing nodes that do many different things: the available nodes are discussed in the next section.

Grouping Compositing Nodes

As your map grows, you might find it helpful to organize your nodes by grouping some of them together. As a suggestion, try to arrange the 'assembly line' of nodes from left to right, top to bottom. You can also group nodes together into subroutines, as discussed below.

Panel: Node Editor → **Node Composition**

Menu: â § ShiftA â Groups

Node maps can get quite complex. Blender allows you to group a set of nodes together to both save space and help you conceptualize what the net effect of a mini-map does to a material/image. This menu selection shows the names of the groups of nodes that you have defined. Select any one to add that group to the map.

Create a Node Group

A group is created by â § Shift-clicking all the nodes you want in the group, and then selecting Node -> Make Group (CtrlG). The group node is shown with a green bar, and the name is shown in an editable field (â § Shift-click on the name to enter EditMode and change the name to something that represents what that group does. The input sockets to the group is the input sockets of contained nodes, and akin for output(s).

Editing Node Groups

You can select the group and press â Tab (or select Node -> Edit Group) to enter/leave the group.

Ungroup Nodes

You can select the group and press AltG (or select Node -> Ungroup) to convert the grouped nodes back to normal.

A "FadeToBlack" compositing group would be a Time node that feeds the Factor of an RBG Curve node. The curve in the RGB Curve node would be a flat line at zero. Namely, any X that came in results in a zero output (black). The factor starts at zero (no effect on the input image), and cranks up to one (full effect, namely black). The input sequence will fade to black over the frame count set in the timer.

Some Examples of Using Compositing Nodes

You can do just about anything with images using nodes. The "Fade to Black" noodle group above is a simple example of what you can do using compositing nodes.

Raw footage from a foreground actor in front of a blue screen, or a rendered object doing something can be layered on top of a background. Composite both together, and you have composited footage.

You can change the mood of an image:

- To make an image 'feel' colder, a blue tinge is added to the pic.
- To convey a flashback or memory, the image may be softened.
- To convey hatred and frustration, add a red tinge or enhance the red. The film 'Sin City' is the most extreme example of this I have ever seen.
- A startling event may be sharpened and contrast enhanced.
- A happy feeling - you guessed it - yellow (equal parts red and green no blue) for bright and sunny.
- Dust and airborne dirt is often added as a cloud texture over the image to give a little more realism.

Terms
Compositing works equally well on a static image or a video sequence. We use the term **image** here to represent either. A **Noodle** is a set of connected nodes.

This section explains the window in general, and its header menu options. It also tells you how to enable nodes for use within Blender.

More information on **where to use nodes** please refer to the following pages:

- Nodes for Materials (types of material nodes).
- Nodes for Composition (types of composition nodes).

To learn more about **how to handle nodes themselves** in general please refer to the following page:

- Doc:Manual/Materials/Nodes/Usage.

Also there is a reference page available explaining the nodes windows and connected functions:

- Doc:Reference/Windows/Nodes.

Accessing The Node Editor



Select the Node Editor window.
First let's enter the node editor by changing our window type to Node Editor. As shown in *Select the Node Editor window*, click on the window type icon and select Node Editor from the popup list. Node maps can get quite large, so use or create a big window. The window has a graph-paper style background and a header.

Each scene within your blend file can have multiple Material Node map and ONE Compositing Node map. The Node Editor window shows either map, depending on the selector position.

Hint
You might want to add a new window layout called 6-Nodes (the list is shown on the User Preferences header at the top of your screen) comprised mostly of one big Node Editor window. My layout has the buttons window at the bottom and a text editor window on the side for me to keep notes. If you have a widescreen display (or even a regular one), you might also want to add a 3D view or UV/Image Editor window to the left side of the Node window layout, so you can work with images or your model while you're manipulating nodes. Having the 3D Preview Render panel open on top of an object is quite useful if you're tweaking material nodes.

By default, the header, when first displayed, is uninitialized as shown:



Default Node Editor header.

Activating Nodes

- What nodes to use?
 - ♦ If you want to work with a material node map, click the ball in the Material/Compositing node set selector. (see *Node Editor Header with Material Nodes enabled.*)
 - ♦ If you want to work with a compositing node map, click the face on the Material/Compositing node set selector. (see *Node Editor Header with Compositing Nodes enabled.*)
- To actually activate nodes, click the Use Nodes button.
- The first time that you select either a Material or a Compsiting node map, the Node Editor window will be instantly filled with starter input and output compositing nodes already connected together.

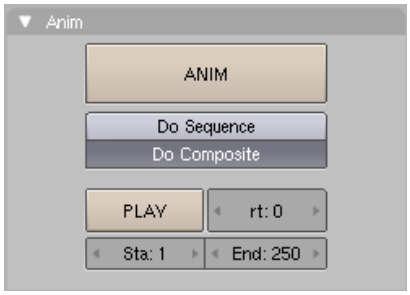


Node Editor Header with Material Nodes enabled.



Node Editor Header with Compositing Nodes enabled.

Required settings for Composition



Use Composition Nodes.
If you are compositing, you must now tell Blender to use the Node map that has been created, and to composite the image using the Node Map. To do so, click on the Do Composite button located below the Animation button. This tells Blender to composite the final image by running it through the composition node map.

From here, you add and connect nodes in a sort of map layout to your heart's content (or physical memory constraints, whichever comes first). But first, let's lay the groundwork by going over the window in general and the header menu options and buttons.

Node Editor Window Actions

When the cursor is in the window, several standard Blender hotkeys and mouse actions are available, including:

Popup menu
: Space - Brings up a main popup menu, allowing you to add, view, select, etc.

Delete
: X or Del - Deletes the selected node(s).

Box select
: B - Starts the bounding box selection process. Position your cursor and LMB click & drag to select a set of nodes.

Cut connections (box)
: LMB click & drag - Starts a box selection, BUT when you let up the mouse button, all threads (connections) within the box are broken.

Undo
: CtrlZ Very helpful if you forgot to press B before box-selecting, eh?

Redo

CtrlY or Ctrlâ § ShiftZ - You can use this if you used "undo" a bit to often :)

Select multiple
â § ShiftLMB or â § ShiftRMB - Multiple node select.

Grab/Move
G - Moves your current selection around.

Execute
E - pumps inputs through the noodle, refreshing everything.

Standard Window Control
Node maps can get pretty hairy (large and complicated, that is). The contents of the window, (the node map) can be panned just like any other Blender window by clicking MMB and dragging about. Wheeling Wheel up/down or using the keypad
+ NumPad/- NumPad will zoom in/out. The window can be resized and combined using the standard window techniques (see *Navigating in 3d Space*).

Node Editor Header



Node Editor Header with Material Nodes enabled.



Node Editor Header with Compositing Nodes enabled.

On the window header, you will see header options:

- View - to see things more clearly;
- Select - to do things more clearly;
- Add - to walk with...err..to add Nodes, organized by type;
- Node - to do things with selected nodes, akin to vertices;
- a Material or Compositing node set selector;
- a Use Nodes button;
- a Free Unused button.

View, Select, and Add Header Menus

These popup menus provide the basic functions:

View
This menu changes your view of the window, standing in for the standard keyboard shortcuts + NumPad (zoom in), - NumPad (zoom out), â Home (zoom all) or equivalent mouse actions.

Select
This menu allows you to select a node or groups of nodes, and does the same as typing the hotkey to select all A or start the border select B process.

Add

This menu allows you to add nodes. Please see the next section for a discussion on the types of nodes that you can add, and what they do. Clicking this menu item is the same as pressing Space when the cursor is in the window

Node Header Menu

Show Cyclic Dependencies

C - Ok, so you've been adding and connecting nodes to your heart's content, and you haven't run out of memory yet. Selecting Show Cyclic Dependencies will show you where you have connected your threads in a circle. For example, you can easily connect a mix output as input to another node, and then connect that node's output back to the mix node input, resulting in a little circle where the image just runs round and round. Left alone, it will eventually get tired and dizzy and crash your computer.

Hide

H - Hides your selected nodes. Just like vertices in a mesh.

Grouping

Most importantly, this menu option allows you to create a user-defined group of nodes. This group can then be edited and added to the map. To create a group, select the nodes you want, and then Node â Make Group, or just use the keyboard shortcut CtrlG. Edit the name using the little input box in the group. Groups are easily identified by their green header and cool names you have picked for them.

Delete

X - Deletes selected nodes.

Duplicate

â § ShiftD - Makes an Unlinked copy, with the same settings as the original.

Grab

G - Moves the little nodes around according to your mouse, just like with meshes.

Duplicate - Faked you out

The new copy is placed **exactly over the old one**. But it isn't the connected one, so playing with the controls will do nothing to your images, even though it **looks** like it's connected with the little threads coming out of the node that is **underneath**. You have to move the duplicated node to reveal the connected node beneath it.

Grab - Reminder Only

Just like my mother-in-law, the menu item does not actually do anything; it's just there to remind you that you can press the G key when your cursor is in the window and actually accomplish something with your life (like rearranging nodes in the window).

Material/Composting Selector Button

Nodes are grouped into two categories, based on what they operate on. Material Nodes operate on a material in use within the blend file. To work with Material Nodes, click on the ball. When you want to work with Compositing nodes, click on the face to show the Compositing Node map.

Use Nodes Header Button

This button tells the render engine to use the node map in computing the material color or rendering the final image, or not. If not, the map is ignored and the basic render of the material tabs or scene is accomplished.

Free Unused Header Button

This button frees up memory space when you have a very complex node map. Recommended.

Node Basics

What are nodes

"Nodes" (some people call connected nodes "noodles") are individual blocks that perform a certain operation. Conceptually, there are three basic types of nodes:

- Nodes that *produce* information, but do not have any inputs of their own. The **Render Layers**,**Value** and **RGB** nodes are examples of this.
- Nodes that *filter* or *transform* their inputs, to produce one or more outputs. The **RGB Curves, Defocus,** and **Vector Blur** nodes are an example.
- Nodes that *consume* their inputs to produce some kind of meaningful result, such as **Composite** (which determines the final output used by Blender), **Viewer** (which displays things right in the noodle display, and **File Output**. A single noodle can produce many different outputs.

The essential idea of nodes is that you can create an arbitrarily-complex *network* of nodes, by connecting the *outputs* of one or more nodes to the *inputs* of one or more other nodes. Then, you can set appropriate parameters (as you see fit) for each node. This network ("noodle") describes how image and texture information literally *flows through* the "noodle" to produce whatever result you want.

As you'll see in just a moment, you can also define *node groups*, then treat those groups as though they were a single node. (Rather like a "macro" or a "user-defined function" in a programming language...) You can link and append these node groups from other files.

Adding and Arranging Nodes

Nodes are added in two ways to the node editor window:

- By clicking the Add menu in the node editor toolbar and picking the type of node you want, or
- By having your cursor in the node editor window and pressing Space and picking a node from the popup Add menu.

In general, try to arrange your nodes within the window such that the image flows from left to right, top to bottom. Move a node by clicking on a benign area and dragging it around. The node can be clicked almost anywhere and dragged about; connections will reshape as a bezier curve as best as possible.

Threads beneath Nodes
Threads (the curves that connect sockets) may reposition behind a node; however they are just that and do not interact with that node in any way.

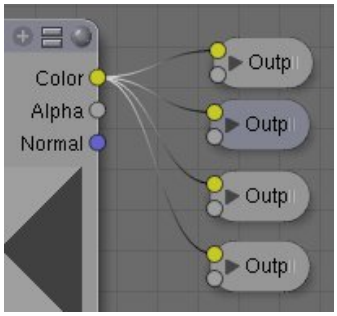
Nodes beneath Nodes
"Gotcha!" When you duplicate a node, the new node will be positioned *exactly* on top of the node that was duplicated. If you leave it there (and it's quite easy to do so), you can't easily tell that there are *two* nodes there! Sure, the drop-shadow may be slightly darker than before, but it's no longer visually obvious just what's connected to what. When in doubt, grab a node and move it slightly to see if something's lurking underneath.

Sockets



Node Sockets.
Each Node in your node window will have "sockets" (often also referred to as "connectors") which are small colored circles to which input data and output data will be linked (*Node Sockets*).

The sockets on the left side of a node describe *inputs*, while the sockets on the right side are *outputs.*



Node Linking.
For your convenience, nodes are *color-coded* according to the type of information they expect to send or receive. There are three colors:

Yellow sockets
: Indicates that **color** information needs to be input or will be output from the node.

Grey sockets
: Indicates values (**numeric**) information. It can either be a single numerical value or a so-called "value map." (You can think if a value map as a grayscale-map where the different amount of bright/dark reflects the value for each point.) If a single value is used as an input for a "value map" socket, all points of the map are set to this same value.
Common use: Alpha maps and value-options for a node.

Blue/Purple sockets
: Indicates **vector/coordinate/normal** information.

Between nodes, yellow must be linked to yellow, gray to gray, blue to blue, unless you use a *converter*, which we'll cover later on.

Next to the color in the node you will see the name of that socket. Though not always the case, you can see the name of the socket as what the information is *intended* to be. But this is not necessarily what it *has* to be. For example, I can add a link from an gray socket titled Alpha to the material node's gray Reflection socket

and still get a result, they key thing being that it's a "gray to gray" connection.

There are exceptions where you can mix yellow (*e.g.* a color-image) and gray (*e.g.* grayscale) without converters. Blender normally places a converter if needed, so feel free to experiment with them. You can use the "Viewer" output nodes, as explained in the later sections, to see if/how it works.

Connecting and Disconnecting Sockets

You link between sockets by clicking the socket with the LMB and holding to drag the thread to another socket, you then let go once you reach the corresponding socket. To break a link between sockets click the LMB and hold and drag a box around any part (it can be really small) to break the link. From output sockets, multiple threads can be extracted and attached to many nodes (*Node Linking*).

In this case, a copy of each output is routed along a thread. However, only a single thread can be linked to an input socket.

Node Groups

Both material and composite nodes can be grouped. Grouping nodes can simplify the node network layout in the node editor, making your material or composite 'noodle' (node network) easier to work with. Grouping nodes also creates what are called NodeGroups (inside a .blend file) or NodeTrees (when appending).

Conceptually, "grouping" allows you to specify a *set* of nodes that you can treat as though it were "just one node." You can then re-use it one or more times in this or some different .blend file(s).

As an example: If you have created a material using nodes that you would like to use in another .blend file, you *could* simply append the material from one .blend file to another. However, what if you would like to create a new material, and use a branch from an existing material node network? You could re-create the branch. Or you could append the material to the new .blend file, then cut and paste the branch that you want into the new material. Both of these options work, but are not very efficient when working across different .blend files. What if you have created a â← Depth of Fieldâ← composite node network and would like to use it in another .blend file? What if you wanted to apply exactly the same series of operations, dozens of times? Here again, you *could* re-create the network, but this is not very efficient. A better method of re-use, for either material node branches or composite node networks, would be to create groups of nodes.

Once a group has been defined, it becomes an opaque object; a reusable software component. You can (if you choose) ignore exactly how it is *defined*, and simply use it (as many times as you like) for whatever it *does*. Groups can be made available through the Blender's library and standard appending method.

Grouping Nodes

To create a node group, in the node editor, select the nodes you want to include, then press CtrlG or Space â node â make group. A node group will have a green title bar. All of the selected nodes will now be minimized and contained within the group node. Default naming for the node groups is *NodeGroup, NodeGroup.001* etc. There is a name field in the node group you can click into to change the name of the group. Change the name of the node group to something meaningful. When appending node groups from one .blend file to another, Blender does not make a distinction between material node groups or composite node groups, so I recommend some naming convention that will allow you to easily distinguish between the two types. For example, name your material node branches *Mat_XXX*, and your composite node networks *Cmp_XXX*.

What NOT to include in your groups.

Remember that the essential idea is that a group should be an easily-reusable, self-contained, software component.

Material node groups should not include material nodes or output nodes.

If you include a material node in your group, you'll end up having the material node appear *twice:* once inside the group, and once outside the group in the new material node-network. If you include an output node *in* the group, there won't be an output socket available *from* the group!

Composite node groups cannot include a *Render Layer* node (Blender won't let you), and should not contain a *Composite* output node. Here again, if they include any connected Output node (*Viewer, Split Viewer,* etc.) then the Group will not have an image output socket.

Editing Node Groups

With a group node selected, pressing Tab expands the node to a window frame, and the individual nodes within it are shown to you. You can move them around, play with their individual controls, re-thread them internally, etc. just like you can if they were a normal part of your editor window. You will not be able to thread them to an outside node directly from them; you have to use the external sockets on the side of the Group node. To add or remove nodes from the group, you you need to ungroup them.

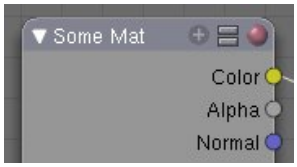
Ungrouping Nodes

The AltG command destroys the group and places the individual nodes into your editor workspace. No internal connections are lost, and now you can thread internal nodes to other nodes in your workspace.

Appending Node Groups

Once you have appended a NodeTree to your .blend file, you can make use of it in the node editor by pressing Space â Add â Groups, then select the appended group. The "control panel" of the Group is the individual controls for the grouped nodes. You can change them by working with the Group node like any other node.

Node Controls



Top of a Node.
At the top of a node there are up to 4 visual controls for the node (*Top of a Node*). Clicking these controls influences how much information the node shows.

 **Arrow**

The arrow on the left collapses the node entirely (*Collapsing Arrow*).


Plus sign (+)

The "plus" icon collapses all sockets that do not have a thread connected to it (*Plus Sign*).


Two squares (=) or "Equal sign"

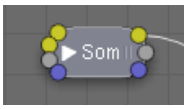
The icon with the two squares collapse all the items in a node that have boxes with information in them (*Menu Collapse*).


Sphere

The sphere icon collapses the viewing window (if the node has one) (*Sphere*).

If the Sphere is **red** this can have 3 reasons:

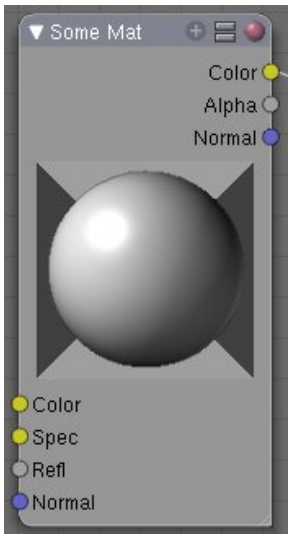
- ◊ It's the only effective output Composite node in the compositor.
- ◊ It's the only effective material Output node (the first one that is added).
- ◊ If it's a Material input node that has a Material (MA:) assigned to it.



Collapsing Arrow.



Plus Sign.



Menu Collapse.

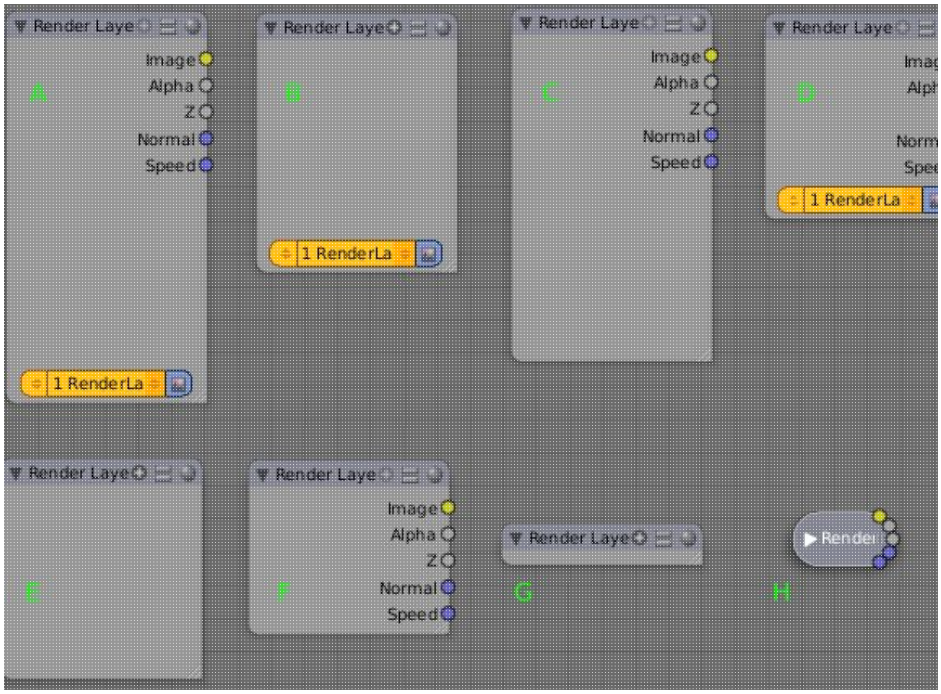


Sphere.



In Combination.

The later three can be used in varying combinations with each other. The arrow that collapses the entire node can only be used in combination with the plus sign (*In Combination*).



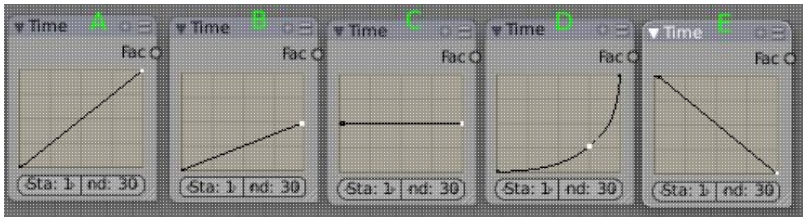
Top sizing controls of a Node
A) Normal, **B**) + Sign clicked, **C**) = Sign clicked, **D**) Sphere clicked, **E**) + and = clicked,
F) = and Sphere clicked, **G**) All three clicked **H**) Arrow clicked.

Node Sizing

Fine Sizing of an individual node can also be accomplished somewhat by clicking LMB and dragging in the lower right-hand corner (where the little slanted lines are).

Node Curves

Some nodes have a curve area that translates an input value to an output value. You can modify this curve shape by clicking on a control point and moving it, or adding a control point. Some examples are shown below:



Modifying a curve node.
Every curve starts out as a straight line with a slope of 1. (My daughter NEVER thought she would use her high school algebra. Ha!) The curve starts out with two tiny black control points at each end of the line. Clicking LMB on a control point selects it and it turns white.

Changing the curve affects how the output is generated. The input, X, usually proceeds linearly (at regular intervals) across the **bottom** axis. Go up until you hit the curve, and then over to the **right** to determine the Y output for that corresponding X. So, for the second example, as X goes from 0 to 1.0 across the bottom, Y

varies from 0.0 to 0.5. In the third, as X goes from 0.0 to 1.0 across the bottom, Y stays constant at 0.5. So, in the picture above, these curves have the following affect on time: **A** don't affect, **B** slow down, **C** stop, **D** accelerate, and **E** reverse time.

The "Curves" widget is a built-in feature in Blender's UI, and can be used anywhere, provided the curve data itself is being delivered to this widget. Currently it is in use in the Node Editor and in the UV Window.

This widget will map an input value horizontally and return the new value as indicated by the height of the curve.

Note: The fact that one of the points on the curve is "white" in each of these screen-shots is *not* significant: it just means that it happened to be the point most-recently selected by your author when preparing this tutorial. What matters here is the shape of *the curve*, not the position (nor the color) of the control-points that were used to define it.

RGB Curves

Multiple curves can be edited in a single widget. The typical use, RGB curves, has "Combined" result or "Color" ("C") as the first curve, and provides curves for the individual R, G, and B components. All four curves are active together, the "C" curve gets evaluated first.

Selecting curve points

- LMB always selects 1 point and deselects the rest.
- Hold â § Shift while clicking to extend the selection or select fewer points.

Editing curves

- LMB click&drag on a point will move points.
- A LMB click on a curve will add a new point.
- Dragging a point exactly on top of another will merge them.
- Holding â § Shift while dragging snaps to grid units.
- CtrlLMB adds a point.
- Use the X icon to remove selected points.

Editing the view

The default view is locked to a 0.0-1.0 area. If clipping is set, which is default, you cannot zoom out or drag the view. Disable clipping with the icon resembling a #.

- LMB click&drag outside of curve moves the view
- Use the + and - icons to zoom in or out.

Special tools

The wrench icon gives a menu with choices to reset a view, to define interpolation of points, or to reset the curve.

The compositing nodes allow you to manipulate an image by routing it through a map of connected nodes. (Such a map is affectionately known as a "noodle.") A starting image is routed through different nodes that do various things to the image, combined with other inputs or put back together, and finally output. Images can be split into their RGB components, combined (mixed) with other inputs, and layered on top of one another. In film, this process is called post-processing and occurs during the post-production (namely, after the raw shots are 'in the can'). At this phase, the raw footage is handed over to an editor who does all sorts of magic to make a good looking movie. Now you can too!

This section is organized by type of nodes, which are grouped based on similar functions:

- Input - Adds something to the node map, such as an image or a value.
- Output - Displays the result in progress as a small image.
- Color - Manipulates the colors of an image.
- Vector - Manipulate the intensities and reflections of an image
- Filters - Process the image to enhance it, working on adjacent pixels.
- Convertors - Separate the image into its component video, or convert formats.
- Mattes - Generating mattes to mask off areas of an image.
- Distortion - Changing the shape of the image.
- Groups - User-defined groups of nodes.

The simplest way to add a node is to put your cursor in a Node Editor window, and then press Space and click on Add. The popup menu will extend to show you these types of nodes. Click on a node type and the popup menu will extend again to show you specific node types available.

Input nodes *produce* information from some source.

For instance, an input could be:

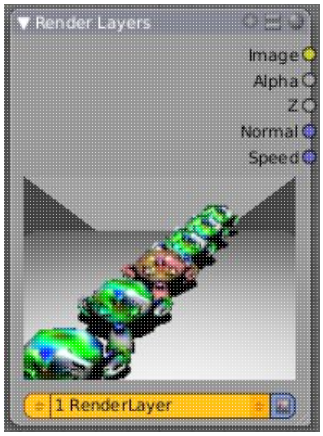
- taken directly from the active camera in a selected scene,
- from a JPG, PNG, etc. file as a static picture,
- a movie clip (such as an animation sequence or home movie), or
- just a color.

These nodes generate the information that feeds other nodes. As such, they have no input-connectors; only outputs.

Render Layers Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Input â Render Layers



Render Layers Node
This node is the starting place to getting a picture of your scene into the compositing node map.

This node inputs an image from a scene within your blend file. Select the scene, and the active render layer from the yellow selection list at the bottom of the node. Blender uses the active camera for that scene to create an image of the objects specified in the RenderLayer.

The Image is input into the map, along with the following data:

- Alpha (transparency) mask

Depending on the Renderlayer passes that are enabled, other sockets are available. By default the Z is enabled:

- Z depth map (how far away each pixel is from the camera)

The example shows that two other passes are enabled:

- Normal vector set (how light bounces off the surface)
- Speed vector set (how fast an object is moving from one frame to the next)

Use the re-render button (Small landscape icon - to the right of the Renderlayer name) to re-render the scene and refresh the image and map.

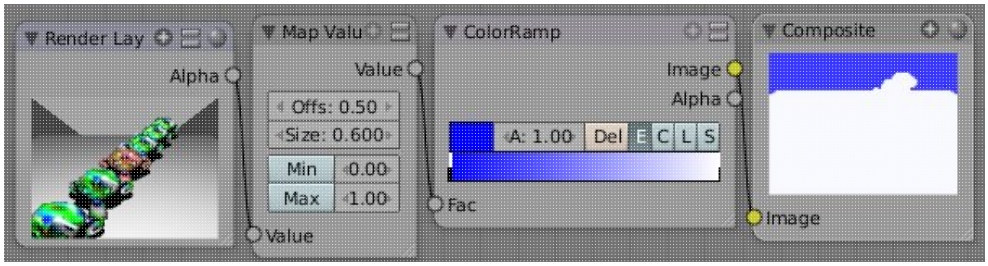
You may recall that a .blend file may contain many scenes. The Renderlayer node can pick up the scene info from any available scene by selecting the scene from the left-hand selector. If that *other* scene also uses the compositor and/or sequencer, you should note that the scene information taken is the raw information (pre-compositing and pre-sequencing). If you wish to use composited information from another scene, you will have to render that scene to a multilayer OpenEXR frameset as an intermediate file store, and then use the Image input node instead.

Using the Alpha Socket

Using the Alpha output socket is crucial in overlaying images on top of one another and letting a background image "show through" the image in front of it.

In a Blender scene, your objects are floating out there in virtual space. While some objects are in front of one another (Z depth), there is no ultimate background. Your world settings can give you the illusion of a horizon, but it's just that; an illusion. Further, some objects are semi-transparent; this is called having an Alpha value. A semi-transparent object allows light (and any background image) to pass through it to the camera. When

you render an image, Blender puts out, in addition to a pretty image, a map of what solid objects actually are there, and where infinity is, and a map of the alpha values for semi-transparent objects. You can see this map by mapping it to a blue screen:



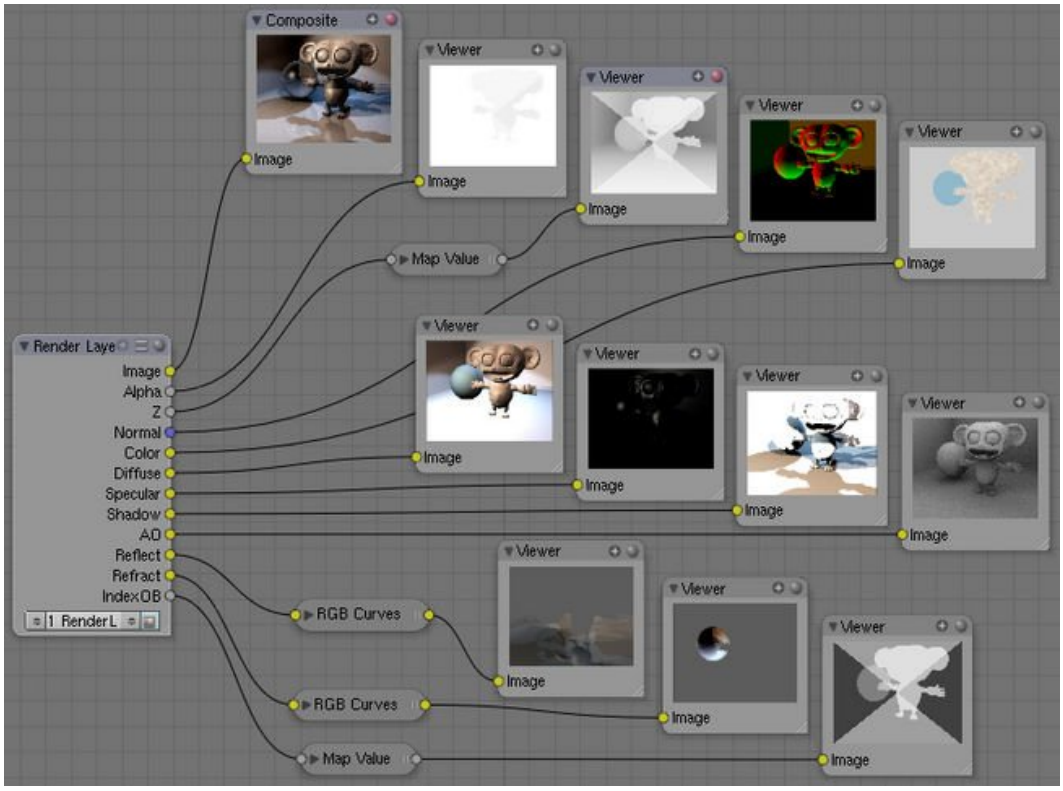
Viewing the Alpha values

In the little node map above, we have connected the Alpha output socket of the RenderLayer node to a Map Value node (explained later, but basically this node takes a set of values and maps them to something we can use). The Color Ramp node (also explained later in detail) takes each value and maps it to a color that we can see with our eyes. Finally, the output of the Color Ramp is output to a Composite viewer to show you, our dear reader, a picture of the Alpha values. Notice that we have set up the map so that things that are perfectly solid (opaque) are white, and things that are perfectly transparent (or where there is nothing) are blue.

Optional Sockets

For any of the optional sockets to appear on the node, you MUST have the corresponding pass enabled. In order for the output socket on the RenderLayer node to show, that pass must be enabled in the RenderLayer panel in the Buttons window. For example, in order to be able to have the Shadow socket show up on the RenderLayer input node, you must have the "Shad" button enabled in the Buttons window, Scene Render buttons, Renderlayer panel. See the RenderLayer tab (Buttons window, Output frame, Render Layers tab, Passes selector buttons) for Blender to put out the values corresponding to the socket.

For a simple scene, a monkey and her bouncy ball, the following picture expertly provides a great example of what each pass looks like:



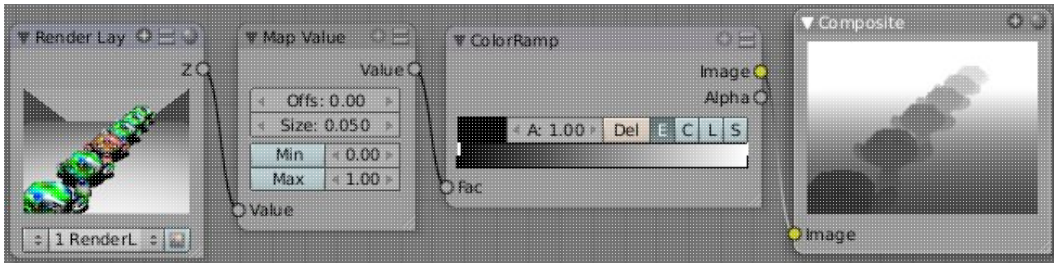
The available sockets are:

- Z: distance away from the camera, in Blender Units
- Normal (Nor): How the color is affected by light coming from the side
- UV: how the image is distorted by the UV mapping
- Speed (Vec): How fast the object is moving, and it what direction
- Color (Col): the RGB values that color the image that you see.
- Diffuse: the softening of colors as they diffuse through the materials
- Specular: the degree of shinyness added to colors as they shine in the light
- Shadow: shadows cast by objects onto other objects
- AO: how the colors are affected by Ambient Occlusion in the world
- Reflect (Ref): for mirror type objects, the colors they reflect and are thus not part of their basic material
- Refract: how colors are bent by passing through transparent objects
- Radio (Radiosity): colors that are emitted by other objects and cast onto the scene
- IndexOB: a numeric ordinal (index) of each object in the scene, as seen by the camera.

Using the Z value Socket

Using the Z output socket is crucial in producing realistic images, since items farther away are blurrier (but more on that later).

Imagine a camera hovering over an X-Y plane. When looking through the camera at the plane, Y is up/down and X is left/right, just like when you are looking at a graph. The camera is up in the air though, so it has a Z value from the X-Y plane, and, from the perspective of the camera, the plane, in fact all the objects that the camera can see, have a Z value as a distance that they are away from it. In addition to the pretty colors of an image, a RenderLayer input node also generates a Z value map. This map is whole bunch of numbers that specify how far away each pixel in the image is away from the camera. You can see this map by translating it into colors, or shades of gray:



Viewing the Z values
In the little node map above, we have connected the Z output socket of the RenderLayer node to a Map Value node (explained later). This node takes a set of values and maps them to something we can use. The Color Ramp node (also explained later in detail) takes each value and maps it to a shade of gray that we can see with our eyes. Finally, the output of the colorramp is output to a Composite viewer to show you, our dear reader, a picture of the Z values. Notice that we have set up the Map Value node so that things closer to the camera appear blacker (think: black is 0, less Z means a smaller number) and pixels/items farther away have an increasing Z distance and therefore get whiter. We chose a Size value of 0.05 to see Z-values ranging from 0 to 20 (20 is 1/0.05).

Using the Speed Socket

Even though things may be animated in our scene, a single image or frame from the animation does not portray any motion; the image from the frame is simply where things are at that particular time. However, from the Render Layers node, Blender puts out a vector set that says how particular pixels are moving, or will move, to the next frame. You use this socket to create a blurring effect. Find out more by clicking here.

Image node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Input â Image

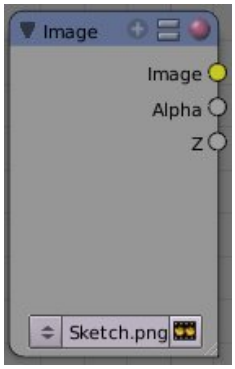


Image node
The Image node injects any image format that is supported by Blender. Besides inputting the actual image, this node can also input Alpha and depth (Z) values if the image has them. If the image is a MultiLayer format, all saved render passes are input. Use this node to input:

- A single image from a file (such as a JPG picture)
- Part or all of an animation sequence (such as the 30th to 60th frame)
- Part or all of a movie clip (such as an AVI file)
- the image that is currently in the UV/Image Editor (and possibly being painted)

- an image that was loaded in the UV/Image Editor

Animated image sequences or video files can also be used. See Animations below.

To select an image file or generated image from the UV/Image Editor, click on the small arrow selector button to the left of the name and pick an existing image (e.g. loaded in the UV editor or elsewhere) or click on LOAD NEW to select a file from your hard disk via a file-browser. These images can be e.g. previously rendered images, matte paintings, a picture of your cat, whatever. Blender really doesn't care.

If the image is part of a sequence, manually click the Image Type selector to the right of the name, and select *Sequence*. Additional controls will allow you to define how much of the sequence to pull in (see Animations below). If the file is a video file, these controls will automatically appear.

Image Channels

When the image is loaded, the available channels will be shown as sockets on the node. As a minimum, the Image, Alpha, and Z channels are made available. The picture may or may not have an alpha (transparency) and/or Z (depth) channel, depending on the format. If the image format does not support A and/or Z, default values are supplied (1.0 for A, 0.0 for Z).

- Alpha/Transparency Channel
 - ♦ If a transparency channel is detected, the Alpha output socket will supply it.
 - ♦ If it does not have an Alpha channel (e.g. JPG images), Blender will supply one, setting the whole image to completely opaque. (An Alpha of 1.00, which will show in a Viewer node as white - if connected to the Image input socket).
- Z/depth Channel
 - ♦ If a Z (depth) channel is detected, the Z output socket will supply it.
 - ♦ If it does not have a Z channel (e.g. JPG or PNG images), Blender will supply one, setting the whole image to be at camera (a depth of 0.00). To view the Z-depth channel, use the Map Value to ColorRamp noodle given above in the Render Layer input node, in the Render Layer input node.

Formats
Blender supports many image formats. Currently only the OpenEXR image format stores RGB (color), A (alpha), and Z (depth) buffer information in a single file, if enabled.

Saving/Retrieving Render Passes



Blender can save the individual Render Layers and specific passes in a MultiLayer file format, which is an extension of the OpenEXR format. In this example, we are reading in frames 50 to 100 of 1 RenderLayer that were generated some time ago. The passes that were saved were the Image, Alpha, Z, Specular and AO pass.

To create a MultiLayer image set, simply disable Do Composite, set your Format to MultiLayer, enable the Render Layer passes you wish to save over the desired frame range, and Animate. Then, in Blender, enable Compositing Nodes and Do Composite, and use the Image input node to read in the EXR file. When you do, you will see each of the saved passes available as sockets for you to use in your compositing noodle.

Image Size

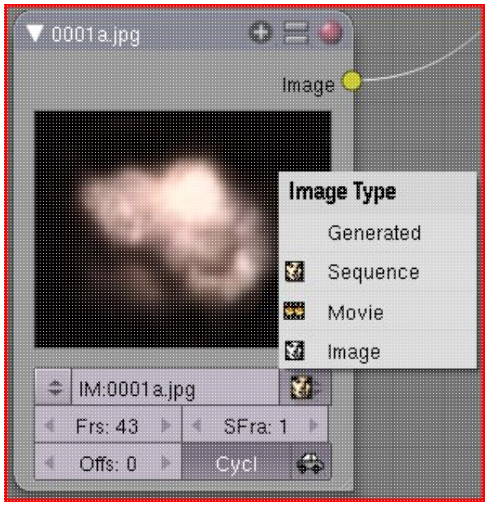
Size matters - Pay attention to image resolution and color depth when mixing and matching images. Aliasing (rough edges), color *flatness*, distorted images, can all be traced to mixing inappropriate resolutions and color depths.

The compositor can mix images with any size, and will only perform operations on pixels where images have an overlap. When nodes receive inputs with differently sized Images, these rules apply:

- The first/top Image input socket defines the output size.
- The composite is centered by default, unless a translation has been assigned to a buffer with the Translate node.

So each node in a composite can operate on different sized images, as defined by its inputs. Only the Composite output node has a fixed size, as defined by the Scene buttons (Format Panel - F10). The Viewer node always shows the size from its input, but when not linked (or linked to a value) it shows a small 320x256 pixel image.

Animations



To use image sequences or movies within your composition, press the face or little film strip button located to the right of the selector. As you click, a pop-up will offer you four choices:

1. Generated -
2. Sequence - a sequence of frames, each frame in a separate file.
3. Movie - a sequence of frames packed into a single .avi or .mov file
4. Image - a single frame or still image in a file

A Movie or Image can be named anything, but a Sequence must have a digit sequence somewhere in its filename, for example fire0001set.jpg, fire0002set.jpg, fire0003set.jpg and so on. The number indicates the frame.

If a Sequence or Movie is selected, an additional set of controls will appear that allows you to select part or all of the sequence. Use these controls to specify which frames, out of the original sequence, that you want to introduce into the animation you are about to render. You can start at the beginning and only use the beginning, or even pick out a set of frames from the middle of an existing animation.

The Frs number button is the number of frames in the sequence that you want to show. For example, if you want to show 2 seconds of the animation, and are running 30 fps, you would put 60 here.

The SFra number button sets the start frame of the animation; namely, at what point in the animation that you *are going to render* do you want this sequence to start playing. For example, if you want to introduce this clip ten seconds into the composite output, you would put 300 here (at 30 fps).

The First number button sets the first number in the animated sequence name. For example, if your images were called "credits-0001.png", "credits-0002.png" through "credits-0300.png" and you wanted to start picking up with frame 20, you'd put 20 here.

To have the movie/sequence start over and repeat when it is done, press the Cyclic button. For example, if you were compositing a fan into a room, and the fan animation lasted 30 frames, the animation would start over at frame 31, 61, 91, and so on, continuously looping. As you scrub from frame to frame, to see the actual video frame used for the current frame of animation, press the auto button to the right of the Cyclic button.

Generated Images



Using the Nodes to modify a painting in progress in the UV/Image window
Blender features Texture Paint which works in the UV/Image Editor, that allows you to paint on the fly, and the image is kept in memory or saved. If sync lock is enabled (the lock icon in the header), changes are broadcast throughout Blender as soon as you lift the mouse button. One of the places that the image can go is to the Image Input node. The example shows a painting session going on in the right-hand UV/Image Editor window for the painting "Untitled". Create this image via Image->New in the UV/Image Editor. Refer to the texture paint section of the user maual for more info on using Texture Paint.

In the left-hand window, the Image input node was used to select that "Untitled" image. Notice that the Image type icon is blank, indicating that it is pulling in a Generated image. That image is colorized by the noodle, with the result used as a backdrop in the Node Editor Window.

Using this setup and the Generated Image type is like painting and post-processing as you continue painting. Changes to either the painting or the post-pro noodle are dynamic and real-time.

Notes

No Frame Stretching or Compression: If the input animation (avi or frame set) was encoded at a frame rate that is *different* from your current settings, the resultant animation will appear to run faster or slower. Blender Nodes do not adjust input video frame rates. Use the scale control inside the Video Sequence Editor to stretch or compress video to the desired speed, and input it here. You can incorporate "Slow-Mo" into your video. To do so, ANIMate a video segment at 60 frames per second, and input it via this node, using Render settings that have an animation frame rate of the normal 30 fps; the resulting video will be played at half speed. Do the opposite to mimic Flash running around at hyperspeed.

AVI (Audio Video Interlaced) files are encoded and often compressed using a routine called a *Codec*. You must have a codec installed on your machine and available to Blender that understands and is able to read the file, in order for Blender to be able to de-code and extract frames from the file. If you get the error message **FFMPEG or unsupported video format** when trying to load the file, you need to get a Codec that understands the video file. Contact the author of the file and find out how it was encoded. An outside package, such as VirtualDub, might help you track this information down. Codecs are supplied by video device manufacturers, Microsoft, DivX, and Xvid, among others, and can often be downloaded from their web sites for free.

Splicing Video Sequences using Nodes

The above animation controls, coupled with a little mixing, is all you need to splice video sequences together. There are many kinds of splices:

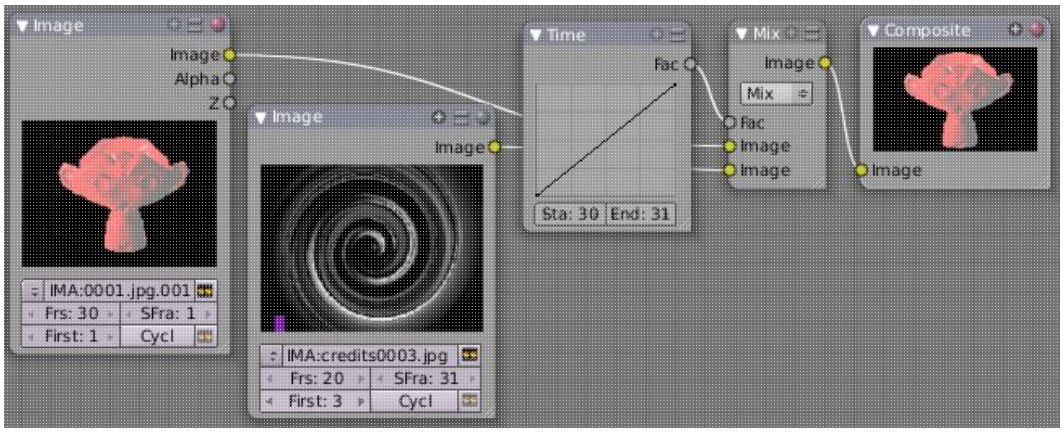
- Cut Splice - literally the ends of the footage are just stuck together
- Fade In - The scene fades in, usually from black
- Fade Out - The scene fades out, usually to black
- Mix - Toward the end of one scene, the images from the next scene meld in as the first scene fades
- Winking and Blinking - fading one cut out while the other fades in, partially or totally through black
- Bumps and Wipes - one cut bumps the other one out of frame, or wipes over it (like from the top left corner down)

Cut Splicing using Nodes

In the example noodle below, we have two pieces of footage that we want to cut splice together.

- Magic Monkey - named 0001.png through 0030.png
- Credits - named credits0001.png through credits0030.png

The editor has reviewed the Credits and thought the first two frames could be thrown away (onto the cutting room floor, as they say) along with the last 8, leaving 20 frames from the total shot. Not shown in this image, but crucial, is that in the Output panel, we set our render output filename to "Monkey-Credits-", and our Animation start and end frames to 1 and 50 (30 from the Monkey, 20 from the credits). Notice the Time node; it tells the Mix node to use the top image until frame 30, and then, at frame 31, changes the Mix factor to 1, which means to use the bottom set of images.

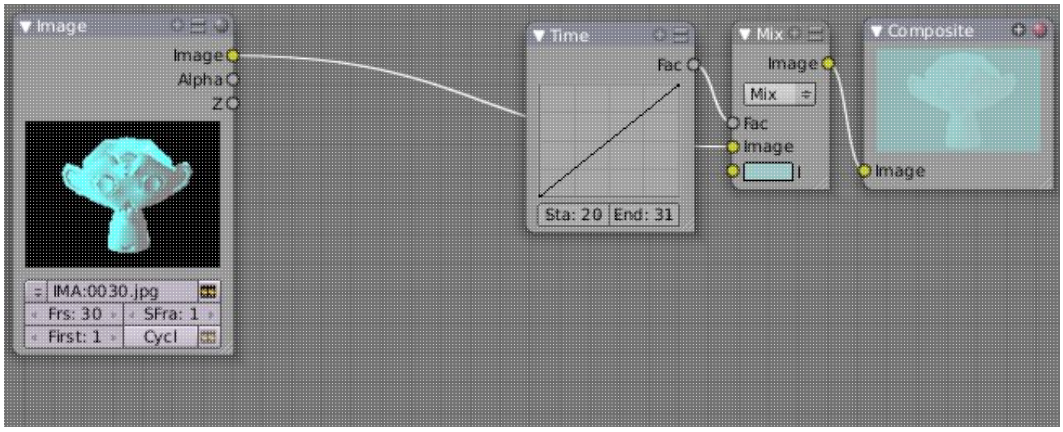


Cut Splice using Nodes
Upon pressing the ANIM button, Blender will composite the animation. If you specified an image format for output, for example, PNG, Blender will create 50 files, named "Monkey-Credits-0001.png" through "Monkey-Credits-0050.png". If you specified a movie format as output, such as AVI-JPEG, then Blender will create only one file, "Monkey-Credits-.avi", containing all 50 frames.

Use cut scenes for rapid-fire transition, conveying a sense of energy and excitement, and to pack in a lot of action in a short time. Try to avoid cutting from a dark scene to a light one, because it's hard on the eyes. It is very emotionally contrasting, and sometimes humorous and ironic, to cut from a very active actor in one scene to a very still actor in another scene, a la old Road Runner and Coyote scenes.

Fade Splicing using Nodes

In the previous topic, we saw how to cut from one sequence to another. To fade in or out, we simply replace one set of images with a flat color, and expand the Time frame for the splice. In the image below, beginning at frame 20, we start fading **out** to cyan:

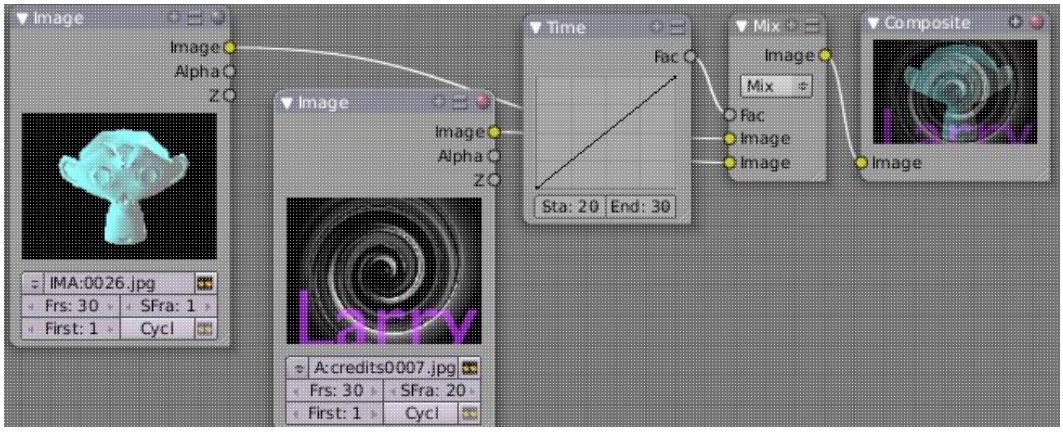


Fading Out using Nodes
Cyan was chosen because that is the color of the Monkey at that time, but you can just as easily choose any color. The image below shows frame 30, when we have almost faded completely.

To fade **in**, change the Mix node and plug the image sequence into the bottom socket, and specify a flat color for the top socket.

Mix Splice using Nodes

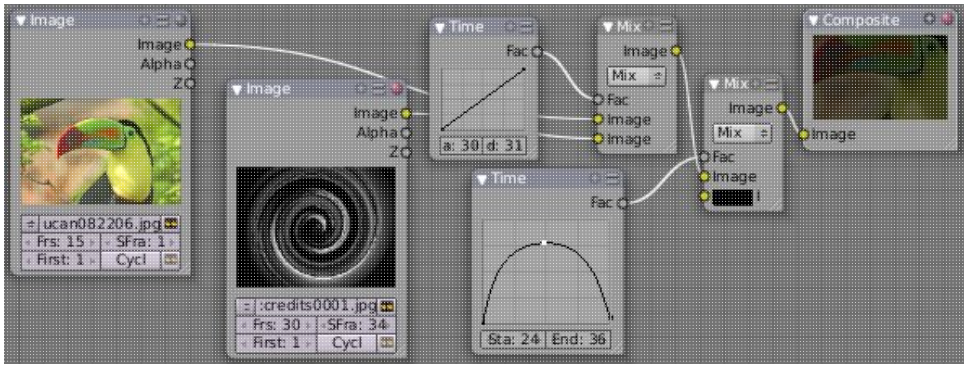
To mix, or crossover, from one scene to the next, start feeding the second scene in while the first is mixing out. The noodle below shows frame 25 of a mix crossover special effect to transition from one scene to the next, beginning at frame 20 with the transition completed by frame 30. Action continues in the first scene as it fades out and is mixed with action that starts in the second scene.



Mix Splice using Nodes
Use this effect to convey similarities between the two scenes. For example, Scene 1 is the robber walking down the street, ending with the camera focusing in on his feet. Scene 2 is a cop walking down the street after him, starting with his feet and working its way up to reveal that the cop is following the robber.

Wink Splice using Nodes

A Wink is just like blinking your eyes; one scene fades to black and the other fades in. To use Blender to get this effect, build on the Cut and Fade splices discussed above to yield:



A Wink using Nodes
In the above example, showing frame 27, we have adjusted some parameters to show you the power of Blender and how to use its Nodes to achieve just the blended crossover effect you desire:

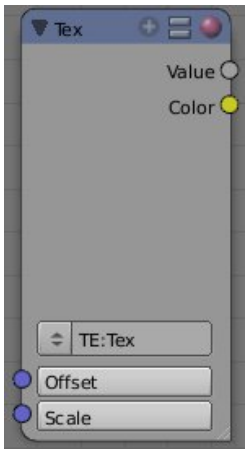
- Postfeed: Even though there were only 15 frames of animation in the Toucan strip, the cutover (top Time node) does not occur until frame 30. Blender continues to put out the last frame of an animation, *automatically extending it for you*, for frames out of the strip's range.
- Prefeed: Even though the swirl does not start playing until frame 34, Blender supplies the first frame of it for Frames 31 through 33. In fact, it supplies this image all the way back to frame 1.
- Partial Fade: Notice the second 'wink' Time node. Like a real wink, it does not totally fade to black; only about 75%. When transitioning between scenes where you want some visual carryover, use this effect because there is not a break in perceptual sequence.

Multiple Feeds
The above examples call out two feeds, but by replicating the Input, Time and Mix nodes, you can have multiple feeds at any one time; just set the Time node to tell the Mixer when to cut over to using it.

Texture Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Input â Texture



Texture node
The Texture node makes 3D textures available to the compositor. A texture, from the list of textures available in the current blend file, is selected and introduced through the value and/or color socket.

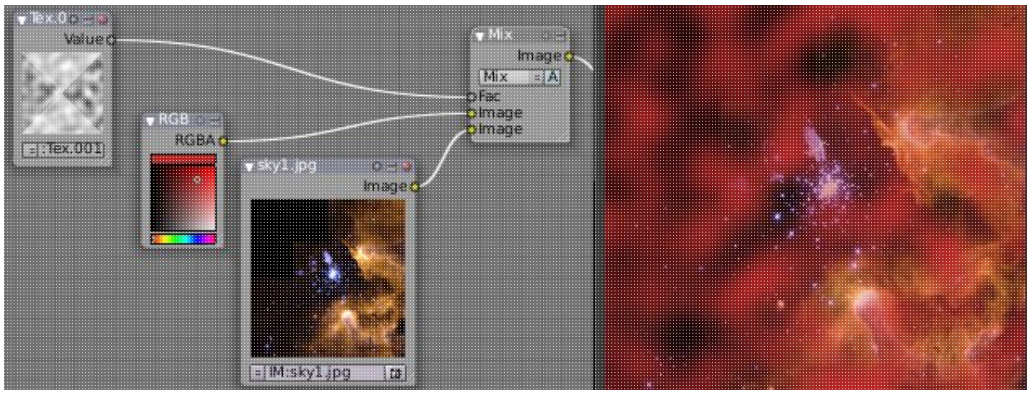
Library
Please read up on the Blender Library system for help on importing and linking to textures in other blender files

Note
You cannot edit the textures themselves in the node window. To use this node, create and edit the texture in the normal texture buttons, then select the texture from the menu button on the node.

You can change the Offset and a Scale (which is called Offs XYZ and Size XYZ in the Materials Texture Map Input panel) for the texture by clicking on the label and setting the sliders, thus affecting how the texture is applied to the image. For animation, note that this is a vector input socket, because the XYZ values are needed.

Texture nodes can output a straight black-and-white Value image (don't mistake this for alpha) and an image (Color).

Example



In the example above, we want to simulate some red plasma gas out there in space. So, we fog up a image taken from the Hubble telecscope of Orion and take the ever-so-useful Cloud texture and use it to mix in red with the image.

Value node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Input â Value



Value Node
The Value node has no inputs; it just outputs a numerical value (floating point spanning 0.00 to 1.00) currently entered in the NumButton displayed in its controls selection.

Use this node to supply a constant, fixed value to other nodes' value or factor input sockets.

RGB node

Panel: Node Editor → Node Composition

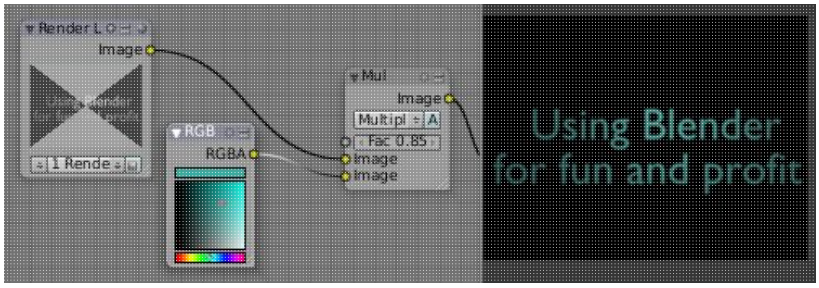
Menu: â § ShiftA â Input â RGB



RGB Node
The RGB node has no inputs. It just outputs the Color currently selected in its controls section; a sample of it is shown in the top box. In the example to the right, a gray color with a tinge of red is slected.

To change the brightness and saturation of the color, LMB click anywhere within the square gradient. The current saturation is shown as a little circle within the gradient. To change the color itself, click anwhere along the rainbow Color Ramp.

Example

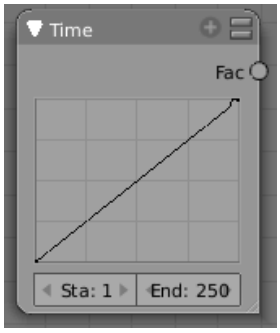


In this example, our corporate color is teal, but the bozo who made the presentation forgot. So, we multiply his lame black and white image with our corporate color to save him from embarassment in front of the boss when he gives his boring presentation.

Time node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Input â Time



Time node
The Time node generates a factor value (from 0.00 to 1.00) (that changes according to the curve drawn) as time progresses through your movie (frames).

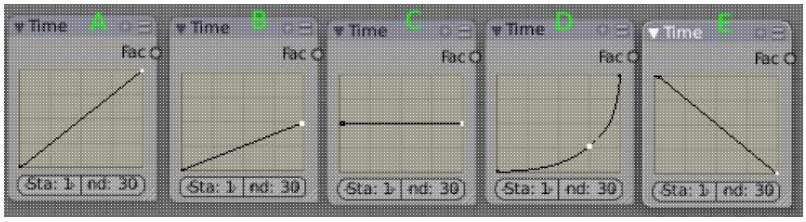
The Start and End NumButtons specify the range of time the values should be output along, and this range becomes the X-axis of the graph. The curve defines the Y-value and hence the factor that is output. In the example to the right, since the timespan is 250 frames and the line is straight from corner to corner, 0.50 would be output at frame 125, and 0.75 will be output at frame 187.

Note on output values
The Map Value node can be used to map the output to a more appropriate value. With some time curves, it is possible that the Time node may output a number larger than one or less than zero. To be save, use the Min/Max clamping function of the Map Value node to limit output.
You can reverse time (unfortunately, only in Blender and not in the real world) by specifying a Start frame greater than the End frame. The net effect of doing so is to flip the curve around. Warning: doing so is easily overlooked in your node map and can be very confusing (like meeting your mother when she was/is your age in "Back to the Future").

Time is Relative
In Blender, time is measured in frames. The actual duration of a time span depends on how fast those frames whiz by (frame rate). You set the frame rate in your animation settings (Scene Context F10). Common settings range from 5 seconds per frame for slideshows (0.2 fps), to 30 fps for US movies.

Time Node Examples

In the picture below, over the course of a second of time (30 frames), the following time controls are made:



A) No Effect **B**) Slow Down **C**) Freeze **D**) Accelerate **E**) Reverse
Common uses for this include a "fade to black", wherein the accelerate time curve (typically exponentially-shaped) feeds a mix value that mixes a constant black color in, so that the blackness accelerates and eventually darkens the image to total black. Other good uses include an increasing soften (blur-out or -in) effect, or fade-in a background or foreground, instead of just jumping things into or out of the scene.

You can even imagine hooking up one blur to a background renderlayer, another inverted blur to a foreground renderlayer, and time feeding both. This node group would simulate someone focusing the camera lens.

Using Input Nodes

As your imagination runs wild, consider a few ideas that came to me just now on my couch: mixing a clouds texture with a time input to fog up a piece of glass or show spray paint building up on a wall. Consider mixing red and the soften with time (decreasing output) to show what someone sees when waking up from a hard hit on the head. Mix HSV input with a starfield image with time (decreasing output) to show what we might see someday as we accelerate our starship and experience red-shift.

As a user, you should know that we have arrived at the point where there are many ways to do the same thing in Blender. For example, an old way to make a slide show using Blender, you created multiple image textures, one image for each slide, and assigned them as texture channels to the material for the screen, then created a screen (plane) that filled the cameral view. Using a material ipo, you would adjust the Color influence of each channel at different frames, fading one in as the previous slide faded out. Whew! Rearranging slide and changing the timing was clunky but doable by moving the IPO keys. The *Node* way is to create an image input, one for each slide image. Using the Image input and Time nodes connected to an AlphaOver mixer is much simpler, clearer, and easier to maintain.

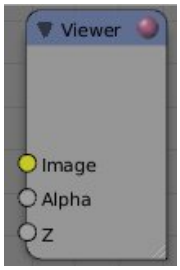
At any point, you may want to see or save the working image in progress, especially right after some operation by a node. Simply create another thread from the image output socket of the node to an Output node to see a mini-picture.

Only one Viewer and one Composite Node is active, which is indicated with a red sphere icon in the Node header. Clicking on Viewer Nodes makes them active. The active Composite Node is always the first, and you should only use one anyway.

Viewer

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Output â Viewer



Viewer node
The Viewer node is a temporary, in-process viewer. Plug it in wherever you would like to see an image or value-map in your node-tree.

LMB click on the image to update it, if it wasn't done automatically. You can use as many of these as you would like.

Using the UV/Image Editor Window

The Viewer node allows results to be displayed in the UV/Image Editor. The image is facilitated by selecting the IM:Viewer Node on the window's header. The UV/Image Editor will display the image from the currently selected viewer node.

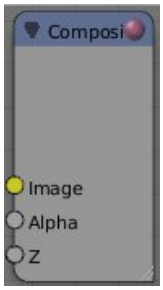
To save the image being viewed, use the Image->Save As menu to save the image in a file.

The UV/Image Editor also has three additional options in its header to view Images with or without Alpha, or to view the Alpha or Z itself. Holding LMB in the Image display allows you to sample the values.

Composite

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Output â Composite



Composite node
The Composite node is where the actual output from the compositor is connected to the renderer. Connecting a node to the Composite node will output the result of that node's full tree to the Renderer, leaving this node unconnected will result in a blank image. This node is updated after each render, but also if you change things in your node-tree (provided at least one finished input node is connected).

You can connect three channels: the actual RGBA image, the Alpha image, and the Z (depth) image. You should only have one Composite node in your map so that only one final image is rendered when the Do Composite button is pressed on the Render Options Animation panel (please read Node Composition if you haven't yet). Otherwise, unpredictable results may occur.

Saving your Composite Image

The RENDER button renders a single frame or image. Save your image using F3 or the File->Save Image menu. The image will be saved using the image format settings on the Render panel.

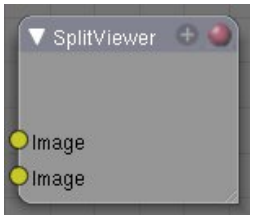
To save a sequence of images, for example if you input a movie clip or used a Time node, with each frame in its own file, use the ANIM button and its settings. If you might want to later overlay them, be sure to use an image format that supports an Alpha channel (such as PNG). If you might want to later arrange them front to back or create a depth of field effect, use a format that supports a Z-depth channel (such as EXR).

To save a composition as a movie clip (all frames in a single file), use an AVI or Quicktime format, and use the ANIM button and its settings.

SplitViewer Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Output â SplitViewer



SplitViewer node
The SplitViewer node takes two images and displays one half of each on each side (top socket on the right half, bottom socket input on the left). Use this node for making side-by-side comparisons of two renderings/images, perhaps from different renderlayers or from different scenes. When transitioning between scenes, you want to be sure the stop action is seamless; use this node to compare the end of one scene with the beginning of another to ensure they align.

File Output Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Output â File Output



File Output node
This node puts out an RGBA image, in the format selected, for each frame range specified, to the filename entered, as part of a frameset sequence. This means that the name of the file will be the name you enter plus a numeric frame number, plus the filename extension (based on format). Based on the format you choose, various quality/compression options may be shown.

To support subsequent arrangement and layering of images, the node can supply a Z-depth map. However, please note that only the OpenEXR image formats save the Z information.

The image is saved whenever Blender feels like it. Just kidding; whenever you press the Render button, the current frame image is saved. When you press the Anim button, the frameset sequence (specified in the Start and End frame) is saved.

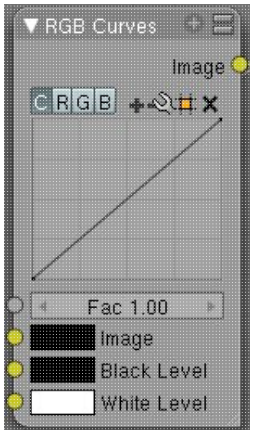
This node saves you from doing (or forgetting to do) the Save Image after a render; the image is saved automagically for you. In addition, since this node can be hooked in anywhere in the noodle, you can save intermediate images automatically. Neat, huh?

Filespecs
As with all filename entries, use // at the beginning of the field to shorthand reference the current directory of the .blend file. You can also use the .. breadcrumb to go up a directory.
These nodes play with the colors in the image. They adjust the image's color intensity, adjust contrast and intensity, and, most importantly, mix two images together by color, transparency, or distance.

RGB Curves Node

Panel: Node Editor → Node Composition

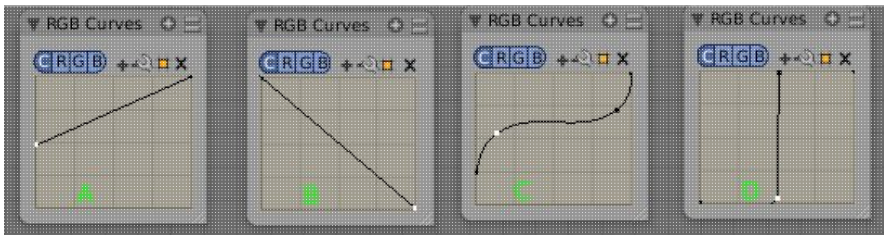
Menu: â § ShiftA â Color â RGB Curves



RGB Curves node
For each color component channel (RGB) or the composite (C), this node allows you to define a bezier curve that varies the input (across the bottom, or x-axis) to produce an output value (the y-axis). By default, it is a straight line with a constant slope, so that .5 along the x-axis results in a .5 y-axis output. Click and drag along the curve to create a control point and to change the curve's shape. Use the X to delete the selected (white) point.

Clicking on each C R G B component displays the curve for that channel. For example, making the composite curve flatter (by clicking and dragging the left-hand point of the curve up) means that a little amount of color will result in a lot more color (a higher Y value). Effectively, this bolsters the faint details while reducing overall contrast. You can also set a curve just for the red, and for example, set the curve so that a little red does not show at all, but a lot of red does.

Here are some common curves you can use to achieve desired effects:



A) Lighten **B**) Negative **C**) Decrease Contrast **D**) Posterize

Options

Fac
: How much the node should factor in its settings and affect the output

Black Level
: Defines the input color that is mapped to black. Default is black, which does not change the image.

White Level
: Defines the input color that is mapped to white. Default is white, which does not change the image.

The levels work exactly like the ones in the image viewer. Input colors are scaled linearly to match black/white levels.

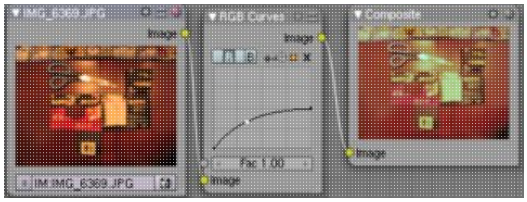
To define the levels, either use LMB on the color patch to bring up the color selection widget or connect some RGBA input to the sockets.

To only affect the value/contrast (not hue) of the output, set the levels to shades of gray. This is equivalent to setting a linear curve for C.

If you set any level to a color with a saturation greater than 0, the output colors will change accordingly, allowing for basic color correction or effects. This is equivalent to setting linear curves for R, G and B.

Examples

Color correction using Curves



Color correction with curves
In this example, the image has way too much red in it, so we run it through an RGB node and reduce the Red channel by about half.

We added a middle dot so we could make the line into a sideways exponential curve. This kind of curve evens out the amount of a color in an image as it reaches saturation. Also, read on for examples of the Darken and Contrast Enhancement curves.

Color correction using Black/White Levels

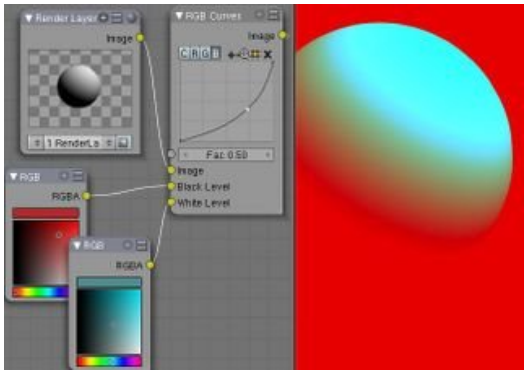


Color correction with Black/White Levels
Manually adjusting the RGB curves for color correction can be difficult. Another option for color correction is to use the Black and White Levels instead, which really might be their main purpose.

In this example, the White Level is set to the color of a bright spot of the sand in the background, and the Black Level to the color in the center of the fish's eye. To do this efficiently it's best to bring up an image viewer window showing the original input image. You can then use the levels' color picker to easily choose the appropriate colors from the input image, zoomig in to pixel level if necessary. The result can be fine-tuned with the R,G, and B curves like in the previous example.

The curve for C is used to compensate the increased contrast that is a side-effect of setting Black and White Levels.

Effects



Changing colors
Curves and Black/White Levels can also be used to completely change the colors of an image.

Note that e.g. setting Black Level to red and White Level to blue does not simply substitute black with red and white with blue as the example image might suggest. Levels do color scaling, not substitution, but depending on the settings they can result in the described color substitution.

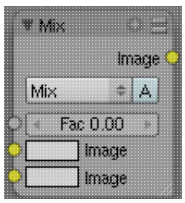
(What really happens when setting Black Level to pure red and White Level to pure blue is that the red channel gets inverted, green gets reduced to zero and blue remains unchanged.)

Because of this the results of setting arbitrary Black/White Levels or RGB curves is hard to predict, but can be fun to play with.

Mix Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Color â Mix



This node mixes a base image (threaded to the top socket) together with a second image (bottom socket) by working on the individual and corresponding pixels in the two images or surfaces. The way the output image is produced is selected in the drop-down menu. The size (output resolution) of the image produced by the mix node is the size of the base image. The alpha and Z channels are mixed as well.

Not one, not two, but count 'em, *sixteen* mixing choices include:

Mix
: The background pixel is covered by the foreground using alpha values.

Add
: The pixels are added together. *Fac* controls how much of the second socket to add in. Gives a bright result.
The "opposite" to Subtract mode.

Subtract
: Pixels are subtracted from one another. Gives a dark result.
The "opposite" to Add mode.

Multiply
: Returns a darker result than either pixel in most cases (except one of them equals white=1). Completely white layers do not change the background at all. Completely black layers give a black result.
The "opposite" to Screen mode.

Screen
: Both pixel values are inverted, multiplied by each other, the result is inverted again. This returns a brighter result than both input pixels in most cases (except one of them equals 0). Completely black layers do not change the background at all (and vice versa) - completely white layers give a white result.
The "opposite" of Multiply mode.

Overlay
: A combination of Screen and Multiply mode, depending on the base color.

Divide
: The background pixel (top socket) is divided by the second one: if this one is white (= 1.0), the first one isn't changed; the darker the second one, the brighter is the result (division by 0.5 - median gray - is same as multiplication by 2.0); if the second is black (= 0.0, zero-division is impossible!), Blender doesn't modify the background pixel.

Difference
: Both pixels are subtracted from one another, the absolute value is taken. So the result shows the distance between both parameters, black stands for equal colors, white for opposite colors (one is black, the other white). The result looks a bit strange in many cases. This mode can be used to invert parts of the base image, and to compare two images (results in black if they are equal).

Darken
: Both pixels are compared to each other, the smaller one is taken. Completely white layers do not change the background at all, and completely black layers give a black result.

Lighten
: Both parameters are compared to each other, the larger one is taken. Completely black layers do not change the image at all and white layers give a white result.

Dodge
: brightens the one socket by the gradient in the other socket. Results in lighter areas of the image where the gradient is whiter. Use the Fac to control how much the gradient affects the other socket.

Burn
: Darkens one socket based on the gradient fed to the other socket. Results in darker images, since the image is *burned* onto the paper, er..image (showing my age).

Color
: Adds a color to a pixel, tinting the overall whole with the color. Use this to increase the tint of an image.

Value
: The RGB values of both pixels are converted to HSV values. The values of both pixels are blended, and the hue and saturation of the base image is combined with the blended value and converted back to RGB.

Saturation
: The RGB values of both pixels are converted to HSV values. The saturation of both pixels are blended, and the hue and value of the base image is combined with the blended saturation and converted back to RGB.

Hue

The RGB values of both pixels are converted to HSV values. The hue of both pixels are blended, and the value and saturation of the base image is combined with the blended hue and converted back to RGB.

Color Channels
There are two ways to express the channels that are combined to result in a color: RGB or HSV. RGB stands for the Red,Green,Blue pixel format, and HSV stands for Hue,Saturation,Value pixel format.

A

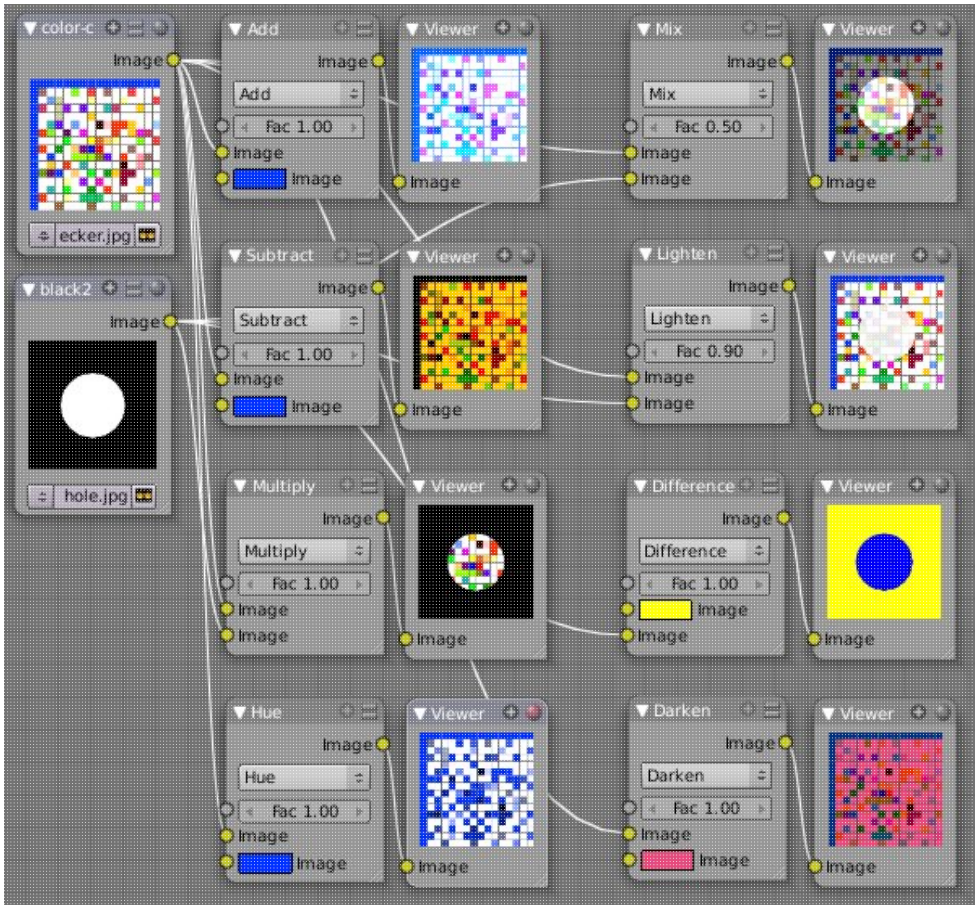
Click the green Alpha button to make the mix node use the Alpha (transparency) values of the second (bottom) node. If enabled, the resulting image will have an Alpha channel that reflects both images' channels. Otherwise, (when not enabled, light green) the output image will mix the colors by considering what effect the Alpha channel has of the base (top input socket) image. The Alpha channel of the output image is not affected.

Fac

The amount of mixing of the bottom socket is selected by the Factor input field (Fac:). A factor of zero does not use the bottom socket, where as a value of 1.0 makes full use. In Mix mode, 50:50 (0.50) is an even mix between the two, but in Add mode, .50 means that only half of the second socket's influence will be applied.

Examples

Below are samples of common mix modes and uses, mixing a color or checker with a mask.

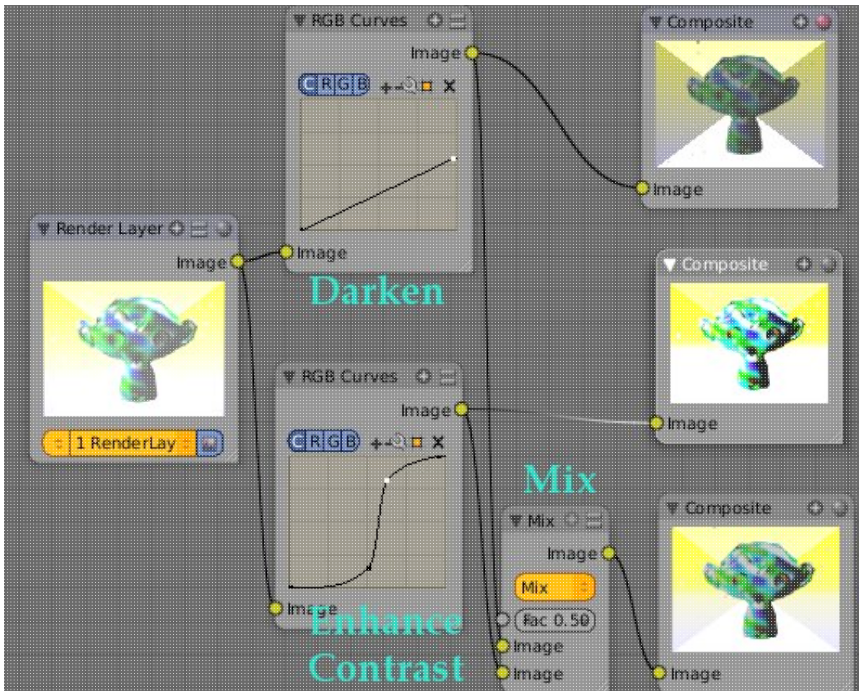


Some explanation of the mixing methods above might help you use the Mix node effectively:

- *Add* - adding blue to blue keeps it blue, but adding blue to red makes purple. White already has a full amount of blue, so it stays white. Use this to shift a color of an image. Adding a blue tinge makes the image feel colder.
- *Subtract* - Taking Blue away from white leaves Red and Green, which combined make Yellow (and you never thought you'd need a color wheel again, eh?). Taking Blue away from Purple leaves Red. Use this to de-saturate an image. Taking away yellow makes an image bluer and more depressing (fun to do, if you're emo).
- *Multiply*: Black (0.00) times anything leaves black. Anything times White (1.00) is itself. Use this to mask out garbage, or to colorize a black-and-white image.
- *Hue*: Shows you how much of a color is in an image, ignoring all colors except what is selected: makes a monochrome picture (style 'Black & Hue').
- *Mix*: Combines the two images, averaging the two
- *Lighten*: Like bleach, makes your whites whiter. Use with a mask to lighten up a little.
- *Difference*: Kinda cute in that it takes out a color. The color needed to turn Yellow into White is Blue. Use this to compare two verrry similar images to see what had been done to one to make it the other; sorta like a change log for images. You can use this to see a watermark you have placed in an image for theft detection.
- *Darken*, with the colors set here, is like looking at the world through rose-colored glasses (sorry, I just couldn't resist).

Contrast Enhancement using Mix

Here is a small map showing the effects of two other common uses for the RGB Curve: **Darken** and **Contrast Enhancement**. You can see the effect each curve has independently, and the combined effect when they are **mixed** equally.



Example node setup showing "Darken", "Enhance Contrast" and "Mix" nodes for composition.
As you can hopefully see, our original magic monkey was overexposed by too much light. To cure an overexposure, you must both darken the image and enhance the contrast. Other paint programs usually provide a slider type of control, but Blender, ah the fantastic Blender, provides a user-definable curve to provide precise control.

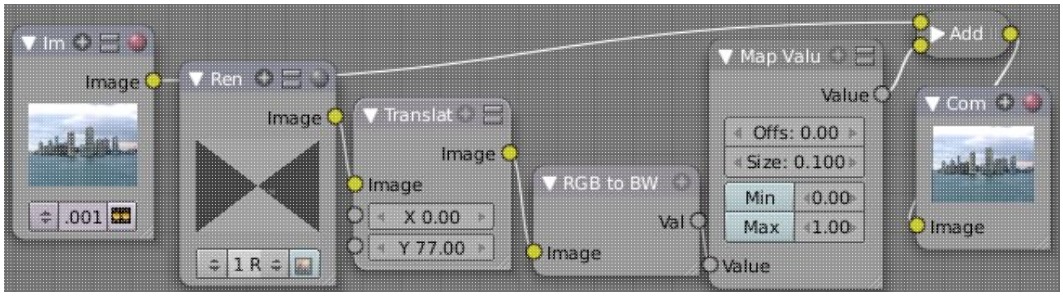
In the top RGB curve, *Darken*, only the right side of the curve was lowered; thus, any X input along the bottom results in a geometrically less Y output. The *Enhance Contrast* RGB 'S' curve scales the output such that middle values of X change dramatically; namely, the middle brightness scale is expanded, and thus whiter whites and blacker blacks are output. To make this curve, simply click on the curve and a new control point is added. Drag the point around to bend the curve as you wish. The Mix node combines these two effects equally, and Suzanne feels much better. And NOBODY wants a cranky monkey on their hands.

Using Mix to Watermark images

In the old days, a pattern was pressed into the paper mush as it dried, creating a mark that identified who made the paper and where it came from. The mark was barely perciptable except in just the right light. Probably the first form of subliminal advertising. Nowadays, people watermark their images to identify them as personal intellectual property, for subliminal advertising of the author or hosting service, or simply to track their image's proliferation throughout the web. Blender provides a complete set of tools for you to both encode your watermark and to tell if an image has your watermark.

Encoding Your Watermark in an Image

First, construct your own personal watermark. You can use your name, a word, a shape or image not easily replicated. While neutral gray works best using the encoding method suggested, you are free to use other colors or patterns. It can be a single pixel or a whole gradient; it's up to you. In the example below, we are encoding the watermark in a specific location in the image using the Translate node; this helps later because we only have to look in a specific location for the mark. We then use the RGB to BW node to convert the image to numbers that the Map Value node can use to make the image subliminal. In this case, it reduces the mark to one-tenth of its original intensity. The Add node adds the corresponding pixels, make the ones containing the mark ever-so-slightly brighter.



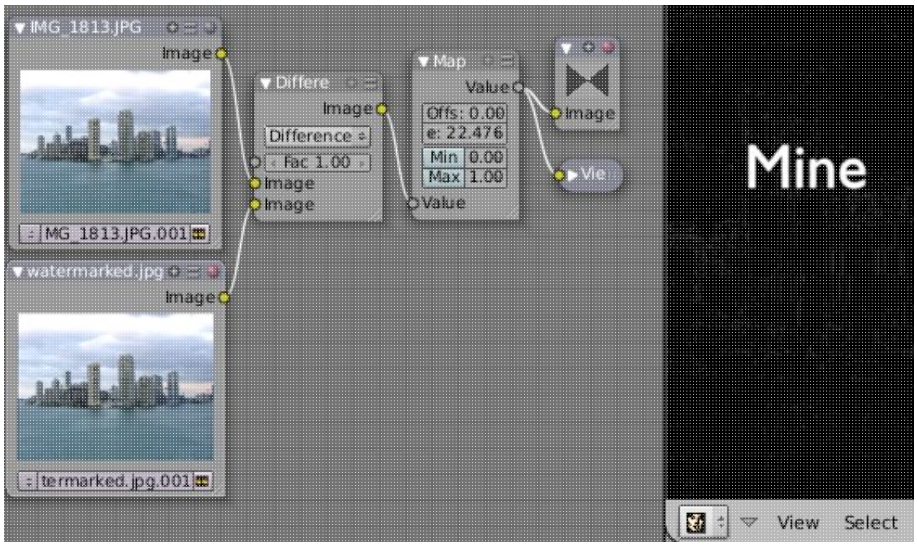
Embedding your mark in an Image using a Mark and Specific Position
Of course, if you *want* people to notice your mark, don't scale it so much, or make it a contrasting color. There are also many other ways, using other mix settings and fancier rigs. Feel free to experiment!

Additional uses
You can also use this technique, using settings that result in visible effects, in title sequences to make the words appear to be cast on the water's surface, or as a special effect to make words appear on the posessed girl's forearm. yuk.

Decoding an Image for your Watermark

When you see an image that you think might be yours, use the node map below to compare it to your stock image (pre-watermarked original). In this map, the Mix node is set to Difference, and the Map Value node amplifies any difference. The result is routed to a viewer, and you can see how the original mark stands out,

clear as a bell:



Checking an image for your watermark
Various image compression algorithms lose some of the original; the difference shows as noise. Experiment with different compression settings and marks to see which works best for you by having the encoding map in one scene, and the decoding map in another. Use them while changing Blender's image format settings, reloading the watermarked image after saving, to get an acceptable result. In the example above, the mark was clearly visible all the way up to JPEG compression of 50%.

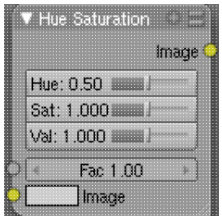
Using Dodge and Burn (History Lesson)

Use the dodge and burn mix methods in combination with a mask to affect only certain areas of the image. In the old darkroom days, when yes, I actually spent hours in a small stinky room bathed in soft red light, I used a circle cut out taped to a straw to dodge areas of the photo as the exposure was made, casting a shadow on the plate and thus limiting the light to a certain area. To do the opposite, I would burn in an image by holding a mask over the image. The mask had a hole in it, letting light through and thus 'burning' in the image onto the paper. The same equivalent can be used here by mixing an alpha mask image with your image using a dodge mixer to lighten an area of your photo. Remember that black is zero (no) effect, and white is one (full) effect. And by the way, ya grew to like the smell of the fixer, and with a little soft music in the background and the sound of the running water, it was very relaxing. I kinda miss those dayz.

Hue Saturation Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Color â Hue Saturation



As an alternative to RGB editing, color can be thought of as a mix of Hues, namely a normalized value along the visible spectrum from infra-red to ultra violet (the rainbow, remember "Roy G. Biv"). The amount of the

color added depends on the saturation of that color; the higher the saturation, the more of that pigment is added. Use the saturation slider of this node to "bring out" the colors of a washed out image.

This node takes an input image and runs the color of the image (and the light it reflects and radiates) 'up' through a factor (0.0-1.0) and applies a saturation of color effect of a hue to the image:

Hue:
: The **Hue** slider specifies how much to shift the hue of the image. Hue 0.5 (in the middle) does not shift the hue or affect the color of the image. As Hue shifts left, the colors shift as more cyan is added; a blue image goes bluer, then greener, then yellow. A red image goes violet, then purple, blue, and finally teal. Shifting right (increasing Hue from 0.5 to 1.0) introduces reds and greens. A blue image goes purple, plum, red, orange, and then yellow. A red image goes golden, olive, green, and cyan.

Sat:
: **Saturation** affect the amount of pigment in the image. A saturation of 0 actually *removes* hues from the color, resulting in a black-and-white grey scale image. A saturation of 1.0 blends in the hue, and 2.0 doubles the amount of pigment and brings out the colors.

Val:
: **Value** affects the overall amount of the color in the image. Increasing value makes an image lighter; decreaing value shifts an image darker.

Fac:
: **Factor** determines how much this node affects the image. A factor of 0 means that the input image is not affected by the Hue and Saturation settings. A factor of 1 means they rule, with .5 being a mix.

Hue Saturation tips

Some things to keep in mind, that might help you use this node better:

Hues are vice-versa.
: A Blue image, with a Hue setting at either end of the spectrum (0 or 1), is output as yellow (Recall that white, minus blue, equals yellow). A Yellow image, with a Hue setting at 0 or 1, is blue.

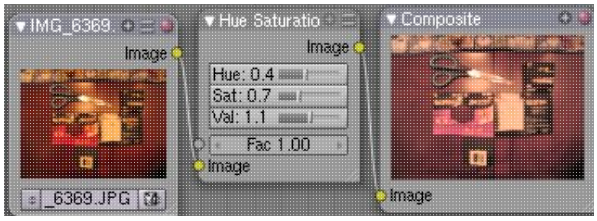
Hue and Saturation work together.
: So, a Hue of .5 keeps the blues the same shade of blue, but the saturation slider can deepen or lighten the intensity of that color.

Gray & White are a neutral hue.
: A grey image, where the RGB values are equal, has no hue. Therefore, this node can only affect it with the Val slider. This applies for all shades of grey, from black to white; wherever the values are equal.

Changing the effect over time.
: The Hue and Saturation values are set in the node by the slider, but you can feed a Time input into the Factor to bring up (or down) the effect change over time.

Tinge
This HSV node simply shifts hues that are already there. To colorize a gray image, or to ADD color to an image, use a mix node to add in a static color from an RGB input node with your image.

HSV Example

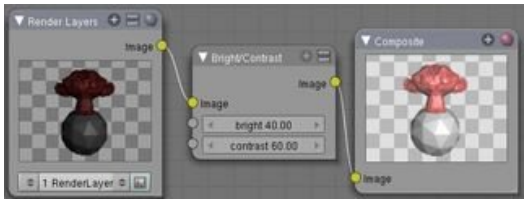


Here, the image taken by a cheap digital camera in poor lighting at night using a flash (can we do it any worse, eh?) is adjusted by decreasing the Hue (decreasing reds and revealing more blues and greens), decreasing Saturation (common in digital cameras, and evens out contrast) and increasing Value (making it all lighter).

Bright/Contrast

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Color â Bright/Contrast

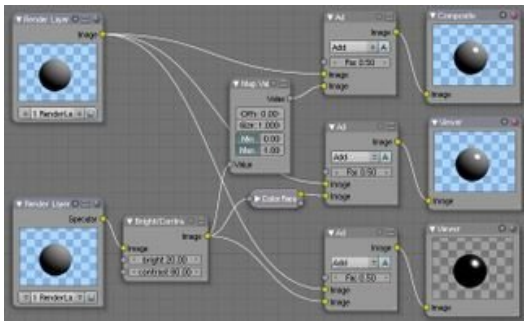


A basic example

Bright
: A multiplier-type factor by which to increase the overall brightness of the image. Use a negative number to darken an image.

Contrast
: A scaling type factor by which to make brighter pixels brighter but keeping the darker pixels dark. Higher values make details stand out. Use a negative number to decrease the overall contrast in the image.

Notes



It is possible that this node will put out a value set that has values beyond normal range, i.e. values > 1 or < 0. If you will be using the output to mix with other images in the normal range, you should clamp the values using the Map Value node (with the Min and Max enabled), or put through a ColorRamp node (with all normal defaults). Either of these nodes will scale the values back to normal range. In the example image, we want to amp up the specular pass. The bottom thread shows what happens if we do not clamp the values; the

specular pass has valued much less than 1 in the dark areas; when added to the medium gray, it makes black. Passing the brightened image through either the Map Value or the ColorRamp produces the desired effect.

Gamma

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Color â Invert



The reason for applying gamma correction to the final render is to correct lighting issues. Lighting issues that are corrected by the gamma correction are light attenuation with distance, light falloff at terminators and light and shadows superpositions. Simply think about the renderer as a virtual camera. By applying a gamma correction to your render, you are just replicating what digital camera do with the photos. Digital camera gamma correct their photos so you do the same thing. The gamma correction is, indeed, 0.45. Not 2.2.

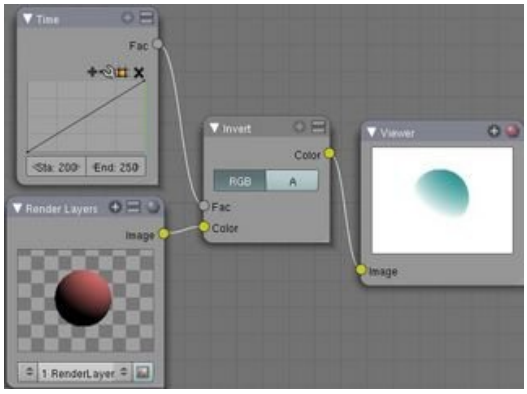
But the reverse gamma correction on the textures and colors have another very important consequence when you are using rendering techniques such as radiosity or GI. When doing the GI calculations, all textures and colors are taken to mean reflectance. If you do not reverse gamma correct your textures and colors, then the GI render will look way too bright because the reflected colors are all way too high and thus a lot more light is bouncing around than it should.

Gamma correction in Blender enters in a few places. The first is in this section with the nodes, both this node and the Tonemap node, and the second is in calculating Radiosity. In the noodle to the left, the split viewer shows the before and after effect of applying a gamma correction.

Invert

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Color â Invert



This handy node inverts the colors in the input image, producing a negative.

Options

Factor
: Controls the amount of influence the node exerts on the output image

Color
: The input image. In this case, a red sphere on a black transparent background

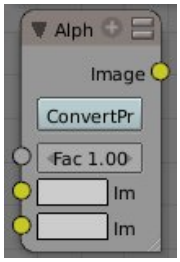
RGB
: Invert the colors from white. In this example, red inverted is cyan (teal).

A
: Invert the alpha (transparency) channel as well. Handy for masking.

AlphaOver Node

Panel: Node Editor → Node Composition

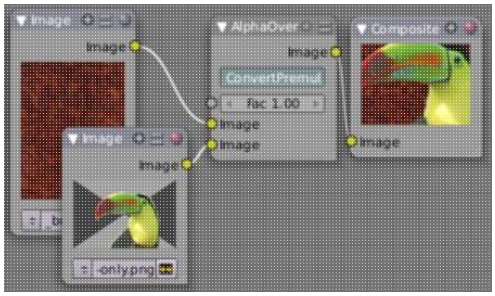
Menu: â § ShiftA â Color â AlphaOver



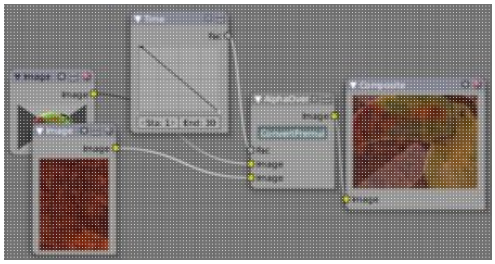
AlphaOver node
Use this node to layer images on top of one another. This node takes two images as input, combines them by a factor, and outputs the image. Connect the Background image to the top input, and the foreground image to the lower input. Where the foreground image pixels have an alpha greater than 0 (namely, have some visibility), the background image will be overlaid.

Use the Factor slider to 'merge' the two pictures. A factor less than 1.00 will make the foreground more transparent, allowing the background to bleed through.

Examples



Assembling a composite Image using AlphaOver
In this example, an image of a Toucan is superimposed over a wooden background. Use the PreMultiply button when the foreground image and background images have a combined Alpha that is greater than 1.00; otherwise you will see an unwanted halo effect. The resulting image is a composite of the two source images.



Animated See-Through/Sheer SFX using AlphaOver - Frame 11
In this example, we use the Factor control to make a sheer cloth or onion-skin effect. You can animate this effect, allowing the observer to 'see-through' walls (or any foreground object) by hooking up a Time node to feed the Factor socket as shown below. In this example, over the course of 30 frames, the Time node makes the AlphaOver node produce a picture that starts with the background wood image, and slowly bleeds through the Toucan. This example shows frame 11 just as the Toucan starts to be revealed.
AlphaOver does not work on the colors of an image, and will not output any image when one of the sockets is unconnnected.

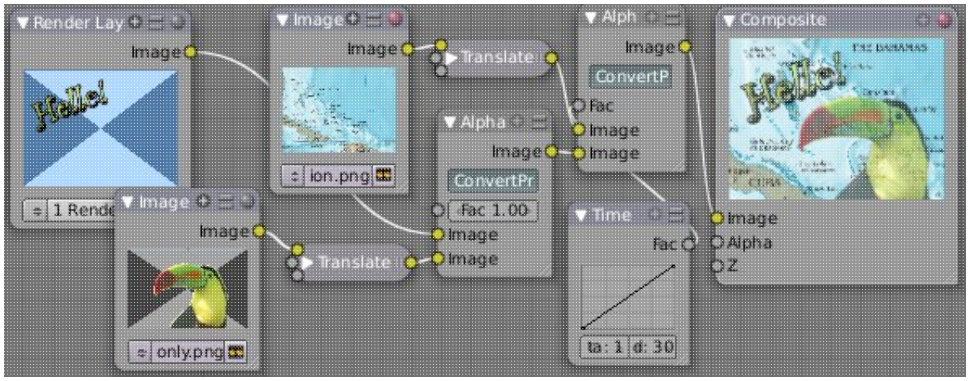
Strange Halos or Outlines

To clarify the premultiplied-alpha button: An alpha channel has a value of between 0 and 1. When you make an image transparent (to composite it over another one), you are really multiplying the RGB pixel values by the alpha values (making the image transparent (0) where the alpha is black (0), and opaque (1) where it is white (1)).

So, to composite image A over image B, you get the alpha of image A and multiply it by image A, thus making the image part of A opaque and the rest transparent. You then inverse the alphas of A and multiply image B by it, thus making image B transparent where A is opaque and vice versa. You then add the resultant images and get the final composite.

A pre-multiplied alpha is when the image (RGB) pixels are already multiplied by the alpha channel, therefore the above compositing op doesn't work too well, and you have to hit 'convert pre-mult'. This is only an issue in semi transparent area, and edges usually. The issue normally occurs in Nodes when you have combined, with alpha, two images, and then wish to combine that image with yet another image. The previously combined image was previously multiplied (pre-mult) and needs to be converted as such (hence, *Convert PreMul*).

If you don't pay attention and multiply twice, you will get a white or clear halo around your image where they meet, since your alpha value is being squared or cubed. It also depends on whether or not you have rendered your image as a pre-mult, or straight RGBA image.

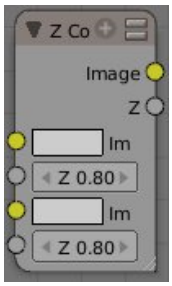


Layering Images using AlphaOver Premul

Z-Combine Node

Panel: Node Editor → Node Composition

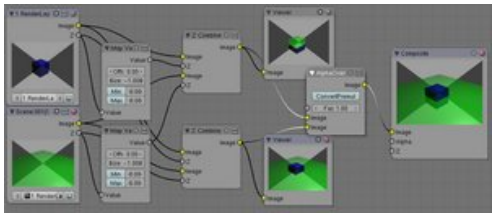
Menu: â § ShiftA â Color â Z Combine



Z Combine node
The Z-Combine node takes two images and two Z-value sets as input. It overlays the images using the provided Z values to detect which parts of one image are in front of the other. If both Z values are equal, it uses the top image. It puts out the combined image, with the combined Z-depth map, allowing you to thread multiple Z-combines together.

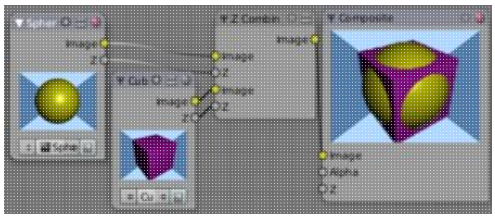
Z-Combine chooses whichever Z-value is less when deciding which image pixel to use. Normally, objects are in front of the camera and have a positive Z value. If one Z-value is negative, and the other positive, Z-Combine will use the image corresponding to the negative value. You can think of a negative Z value as being behind the camera. When choosing between two negative Z-values, Z-Combine will use whichever is more negative.

Alpha values carry over from the input images. Not only is the image pixel chosen, but also its alpha channel value. So, if a pixel is partially or totally transparent, the result of the Z-Combine will also be partially transparent; in which case the background image will show through the foreground (chosen) pixel. Where there are sharp edges or contrast, the alpha map will automatically be anti-aliased to smooth out any artifacts. Hoever, you can obtain this by making an AlphaOver of two Z-Combine, one normal, the other having inverted (reversed?) Z-values as inputs, obtained using for each of them a MapValue node with a Size field set to -1.0:



Alpha and Z-Combine node.

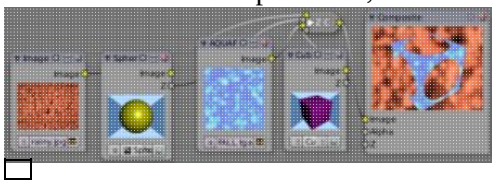
Examples



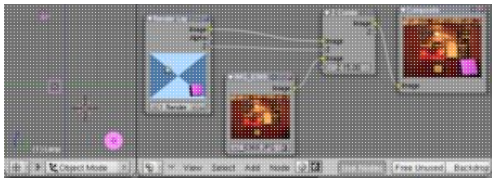
Choosing closest pixels
In the example to the right, render output from two scenes are mixed using the Z-Offset node, one from a sphere of size 1.30, and the other a cube of size 1.00. The sphere and square are located at the same place. The cube is tipped forward, so the corner in the center is closer to the camera than the sphere surface; so Z-Offset chooses to use the cube's pixels. But the sphere is slightly larger (a size of 1.30 versus 1.00), so it does not fit totally 'inside' the cube. At some point, as the cube's sides recede back away from the camera, the sphere's sides are closer. When this happens, Z-offset uses the sphere's pixels to form the resulting picture.

This node can be used to combine a foreground with a background matte painting. Walt Disney pioneered the use of multi-plane mattes, where three or four partial mattes were painted on glass and placed on the left and right at different Z positions; mininal camera moves to the right created the illusion of depth as Bambi moved through the forest.

Valid Input
Z Input Sockets do not accept fixed values; they must get a vector set (see Map Value node). Image Input Sockets will not accept a color, since it does not have UV coordinates.



Mix and Match Images
You can use Z-Combine to merge two images as well, using the Z-values put out by two renderlayers. Using the Z-values from the sphere and cube scenes above, but threading different images, yields the example to the right.



Z-Combine in action
In this noodle (you may click the little expand-o-matic icon in the bottom right to view it to full size), we mix a render scene with a flat image. In the side view of the scene, the purple cube is 10 units away from camera,

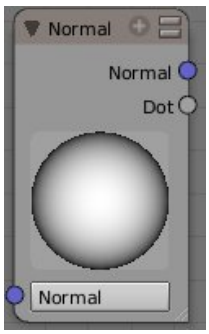
and the gray ball is 20. The 3D cursor is about 15 units away from camera. We Z-in the image at a location of 15, thus inserting it in-between the cube and the ball. The resulting image appears to have the cube on the table.

Invisible Man Effect
If you choose a foreground image which has a higher Alpha than the background, and then mix the Z-combine with a slightly magnified background, the outline of the transparent area will distort the background, enough to make it look like you are seeing part of the background through an invisible yet Fresnel-lens object.
Vector nodes manipulate information about how light interacts with the scene, multiplying vector sets, and other wonderful things that normal humans barely comprehend (except math geniuses, who may not be considered 'normal'). Even if you aren't a math wiz, you'll find these nodes to be very useful.

Normal Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Vector â Normal



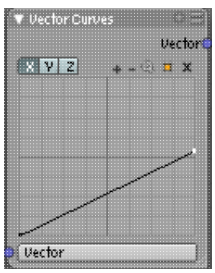
Normal node
The Normal node generates a normal vector and a dot product. Click and Drag on the sphere to set the direction of the normal.

This node can be used to input a new normal vector into the mix. For example, use this node as an input to a Color Mix node. Use an Image input as the other input to the Mixer. The resulting colorized output can be easily varied by moving the light source (click and dragging the sphere).

Vector Curves Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Vector Ops â Vector Curves



The Vector Curves node maps an input vector image's x, y, and z components to a diagonal curve. The three channels are accessed via the X, Y, and Z buttons at the top of the node. Add points to the curve by clicking on it.

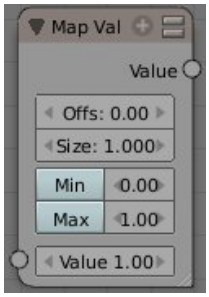
Note that dragging a point across another will switch the order of the two points (e.g. if point A is dragged acros point B, then point B will become point A and point A will become point B).

Use this curve to slow things down or speed them up from the original scene.

Map Value Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Vector â Map Value



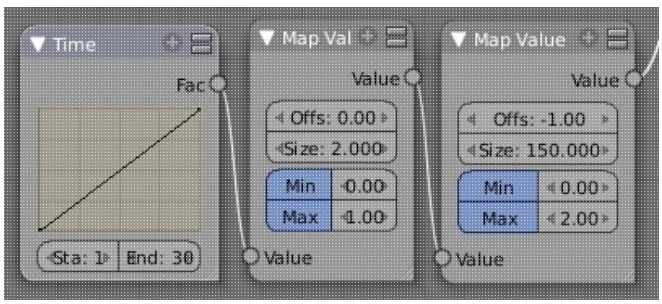
Map Value node
Map Value node is used to scale, offset and clamp values (value refers to each vector in the set). The formula for how this node works is:

- Offs will add a number to the input value
- Size will scale (multiply) that value by a number
- By clicking Min/Max you can set the minimum and maximum numbers to clamp (cut off) the value too. Min and Max must be individually enabled by LMB clicking on the label for them to clamp. â § ShiftLMB on the value to change it.

 ◊ If Min is enabled and the value is less than Min, set the ouput value to Min
 ◊ If Max is enabled and the input value is greater than Max, set the ouput value to Max

This is particularly useful in achieving a depth-of-field effect, where you can use the Map Value node to map a Z value (which can be 20 or 30 or even 500 depending on the scene) to to range between 0-1, suitable for connecting to a Blur node.

Using Map Value to Multiply values



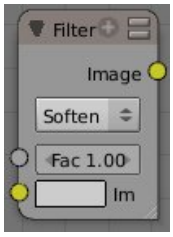
You can also use the map value node to multiply values to achieve an output number that you desire. In the mini-map to the right, the Time node ouputs a value between 0.00 and 1.00 evenly scaled over 30 frames. The *first* Map Value node multiplies the input by 2, resulting in an output value that scales from 0.00 to 2.00 over 30 frames. The *second* Map Value node subtracts 1 from the input, giving working values between -1.00 and 1.00, and multiplies that by 150, resulting in an output value between -150 and 150 over a 30-frame sequence.

Filters process the pixels of an image to highlight additional details or perform some sort of post-processing effect on the image.

Filter Node

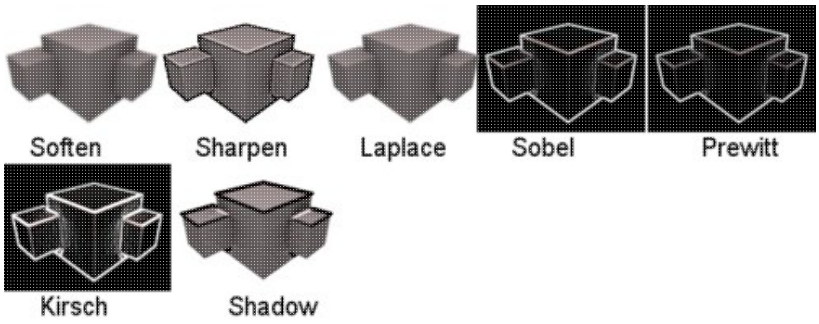
Panel: Node Editor → Node Composition

Menu: â § ShiftA â Filters â Filter



Filter node
The Filter node implements various common image enhancement filters. The supported filters are, if not obvious, named after the mathematical genius who came up with them:

- Soften - Slightly blurs the image.
- Sharpen - Increases the contrast, especially at edges
- Laplace - Softens around edges
- Sobel - Creates a negative image that highlights edges
- Prewitt - Tries to do Sobel one better.
- Kirsch - Improves on the work done by those other two flunkies, giving a better blending as you approach an edge.
- Shadow - Performs a relief emboss/bumpmap effect, darkening outside edges.



The Filter node has seven modes, shown here.
The Soften, Laplace, Sobel, Prewitt and Kirsch all perform edge-detection (in slightly different ways) based on vector calculus and set theory equations that would fill six blackboards with gobbledy gook.
Recommended reading for insomniacs.

Blur Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Filters â Blur

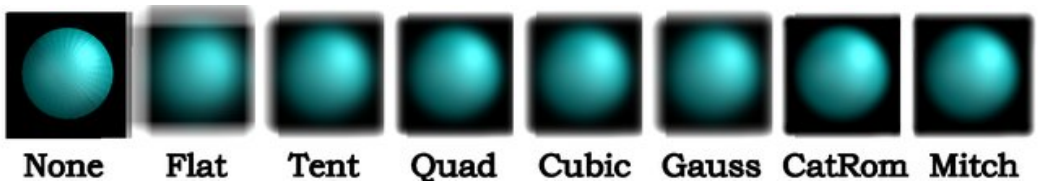


Blur node
The Blur node blurs an image, using one of seven blur modes (set using the upper-left popup button), and a radius defined by the X and Y number buttons. By default these are set to zero, so to enable the node you must set one or both to a value greater then 0. You can optionally connect a value image to the Size input node, to control the blur radius with a mask. The values must be mapped between 0-1 for best effect, as they will be multiplied with the X and Y number button values.

Options

The X and Y values are the number of pixels over which to spread the blur effect.

The Bokeh button (only visible as Bok or Bo on some screen setups) will force the blur node to use a circular blur filter. This gives higher quality results, but is slower then using a normal filter. The Gam button (for "gamma") makes the Blur node gamma-correct the image before blurring it.



Blur node blur modes using 15% of image size as XY, no Bokeh/Gamma. Click expand to see details
The difference between them is how they handle sharp edges and smooth gradients and preserve the highs and the lows. In particular (and you may have to closely examine the full-resolution picture to see this):

Flat just blurs everything uniformly
Tent preserves the high and the lows better making a linear falloff
Quadratic and CatRom keep sharp-contrast edges crisp
Cubic and Mitch preserve the highs but give almost a out-of-focus blur while smoothing sharp edges

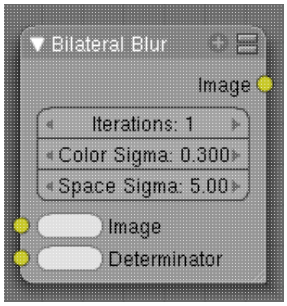
Example

An example blend file, in fact the one used to create the image above, is available here. The .blend file takes one image from the RenderLayer "Blurs" and blurs it while offsetting it (Translate) and then combining it (AlphaOver) to build up the progressive sequence of blurs. Play with the Value and Multiply nodes to change the amount of blurring that each algorithm does.

Bilateral Blur Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Filters â Bilateral Blur



Blur node
The bilateral blur node performs a high quality adaptive blur on the source image. It can be used for various purposes like:

smoothing results from blenders raytraced ambient occlusion
smoothing results from various unbiased renderers,
to fake some performance-heavy processes, like blurry refractions/reflections, soft shadows,
to make non-photorealistic compositing effects.

Inputs

Bilateral blur has 2 inputs:

Image, for the image to be blurred.
Determinator, which is non-obligatory, and is used only if connected.

if only 1st input is connected, the node blurs the image depending on the edges present in the source image. If the Determinator is connected, it serves as the source for defining edges/borders for the blur in the image. This has great advantage in case the source image is too noisy, but normals in combination with zbuffer can still define exact borders/edges of objects.

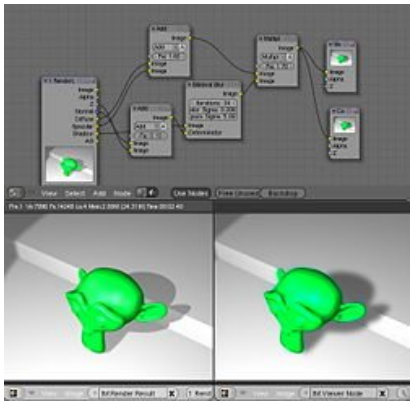
Options

Iterations - defines how many times should the filter perform the operation on the image. It practically defines the radius of blur.

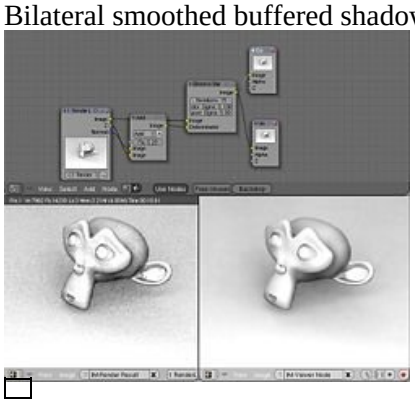
Color Sigma - defines the threshold for which color differences in the image should be taken es edges.

Space sigma - is a fine-tuning variable for blur radius.

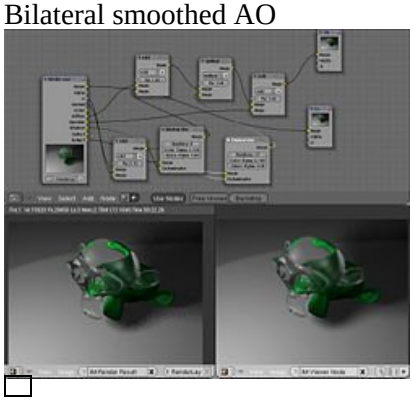
Examples



Bilateral smoothed buffered shadow



Bilateral smoothed AO

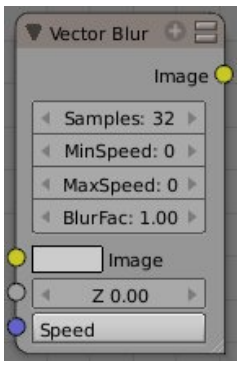


Bilateral faked blurry refraction+smoothed reytraced soft shadow

Vector (Motion) Blur Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Filters â Vector Blur



Vector Blur node
Motion blur is the effect of objects moving so fast they blur. Because CG animations work by rendering individual frames, they have no real knowledge of what was where in the last frame, and where it is now.

In Blender, there are two ways to produce motion blur. The first method (which produces the most correct results) works by rendering a single frame up to 16 times with slight time offsets, then accumlating these images together; this is called Motion Blur and is activated on the Render panel. The second (and much faster) method is the Compositor node Vector Blur.

To use, connect the appropriate passes from a Render Result node.

Note
Make sure to enable the Speed (called Vec) pass in the Render Layers panel for the render layer you wish to perform motion blur on.
Maximum Speed: Because of the way vector blur works, it can produce streaks, lines and other artifacts. These mostly come from pixels moving too fast; to combat these problems, the filter has minimum and maximum speed settings, which can be used to limit which pixels get blurred (e.g. if a pixel is moving really, really fast but you have maximum speed set to a moderate amount, it won't get blurred).

Minimum Speed: Especially when the camera itself moves, the mask created by the vectorblur node can become the entire image. A very simple solution is to introduce a small threshold for moving pixels, which can efficiently separate the hardly-moving pixels from the moving ones, and thus create nice looking masks. You can find this new option as 'min speed'. This minimum speed is in pixel units. A value of just 3 will already clearly separate the background from foreground.

Hint
You can make vector blur results a little smoother by passing the Speed pass through a blur node (but note that this can make strange results, so it's only really appropriate for still images with lots of motion blur).

Examples

An in-depth look at how to use the Vector Blur node can be found here.

As far as we know, this node represents a new approach to calculating motion blur. Use vector blur in compositing with confidence instead of motion blur. In face, when compositing images, it is necessary to use vector blur since there isn't "real" motion. In this example blend file, you will find a rigged hand reaching down to pick up a ball. Based on how the hand is moving (those vectors), the image is blurred in that direction. The fingers closest to the camera (the least Z value) are blurred more, and those farther away (the forearm) is blurred the least.

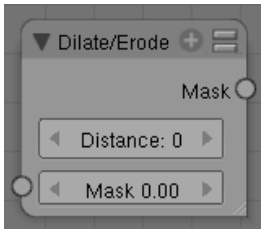
Known Bugs

Blender 2.44+
Does not work when reading from a multilayer OpenEXR sequence set

Dilate/Erode Node

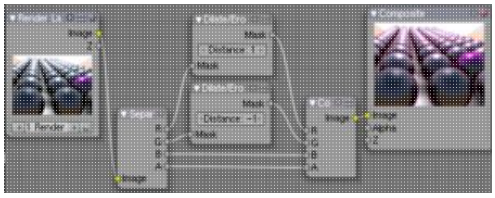
Panel: Node Editor → Node Composition

Menu: â § ShiftA â Filters â Dilate/Erode



Dilate/Erode node
This node blurs individual color channels. The color channel (or a black and white image) is connected to the Mask input socket, and the Distance is set manually (by clicking on the arrows or the value) or automatically from a value node or a time-and-map-value noodle. A positive value of Distance expands the influence of a pixel on its surrounding pixels, thus blurring that color outward. A negative value erodes its influence, thus increases the constrast of that pixel relative to its surrounding pixels, thus sharpening it relative to surrounding pixels of the same color.

Example



Magenta tinge
In the above example image, we wanted to take the rather boring array of ball bearings and spruce it up; make it hot, baby. So, we dilated the red and eroded the green, leaving the blue alone. If we had dilated both red and green...(hint: red and green make yellow). The amount of influence is increased by increasing the Distance values. [Blend file available here.]

Defocus

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Filters â Defocus

This single node can be used to emulate depth of field using a postprocessing method. It can also be used to blur the image in other ways, not necessarily based on 'depth' by connecting something other than a Zbuffer. In essence, this node blurs areas of an image based on the input zbuffer map/mask.

Camera Settings

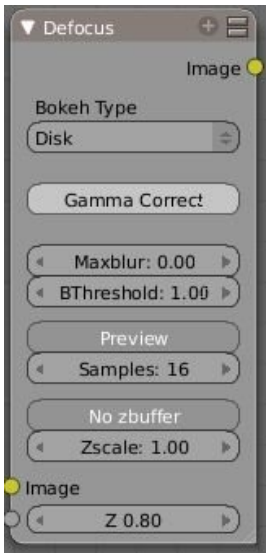


DofDist setting for the camera.
The Defocus node uses the actual camera data in your scene if supplied by a RenderLayer node.

To set the point of focus, the camera now has a DoFDist parameter, which is shorthand for Depth of Field Distance. Use this camera parameter to set the focal plane of the camera (objects DoFDist away from the camera are in focus). Set DofDist in the main Camera edit panel; the button is right below the Lens value button.

To make the focal point visible, enable the camera Limits option, the focal point is then visible as a yellow cross along the view direction of the camera.

Node Inputs



Defocus node
The node requires two inputs, an image and a zbuffer, the latter does not need to be an actual zbuffer, but can also be another (grayscale) image used as mask, or a single value input, for instance from a time node, to vary the effect over time.

Node Setting

The settings for this node are:

Bokeh Type menu
: Here you set the number of iris blades of the virtual camera's diaphragm. It can be set to emulate a perfect circle (Disk) or it can be set to have 3 (Triangle), 4 (Square), 5 (Pentagon), 6 (Hexagon), 7 (Heptagon) or 8 blades (Octagon). The reason it does not go any higher than 8 is that from that point on the result tends to be indistinguishable from a Disk shape anyway.

Rotate
: This button is not visible if the Bokeh Type is set to Disk. It can be used to add an additional rotation offset to the Bokeh shape. The value is the angle in degrees.

Gamma Correct
: Exactly the same as the Gamma option in Blender's general Blur node (see Blur Node). It can be useful to further brighten out of focus parts in the image, accentuating the Bokeh effect.



Defocus node using Z-Buffer

fStop
: This is the most important parameter to control the amount of focal blur: it simulates the aperture *f* of a real lens(' iris) - without modifying the luminosity of the picture, however! As in a real camera, the *smaller* this number is, the more-open the lens iris is, and the *shallower* the depth-of-field will be. The default value 128 is assumed to be infinity: everything is in perfect focus. Half the value will double the amount of blur. This button is not available if No zbuffer is enabled.

Maxblur
: Use this to limit the amount of blur of the most out of focus parts of the image. The value is the maximum blur radius allowed. This can be useful since the actual blur process can sometimes be very slow. (The more blur, the slower it gets.) So, setting this value can help bring down processing times, like for instance when the world background is visible, which in general tends to be the point of maximum blur (not always true, objects very close to the lens might be blurred even more). The default value of 0 means there is no limit to the maximum blur amount.

BThreshold
: The defocus node is not perfect: some artifacts may occur. One such example is in-focus objects against a blurred background, which have a tendency to bleed into the edges of the sharp object. The worst-case scenario is an object in-focus against the very distant world background: the differences in

distance are very large and the result can look quite bad. The node tries to prevent this from occurring by testing that the blur difference between pixels is not too large, the value set here controls how large that blur difference may be to consider it 'safe.' This is all probably quite confusing, and fortunately, in general, there is no need to change the default setting of 1. Only try changing it if you experience problems around any in-focus object.

Preview
: As already mentioned, processing can take a long time. So to help make editing parameters somewhat 'interactive', there is a preview mode which you can enable with this button. Preview mode will render the result using a limited amount of (quasi)random samples, which is a *lot* faster than the 'perfect' mode used otherwise. The sampling mode also tends to produce grainy, noisy pictures (though the more samples you use, the less noisy the result). This option is on by default. Play around with the other parameters until you are happy with the results, and only then disable the preview mode for the final render.

Samples
: Only visible when Preview is set. Sets the amount of samples to use to sample the image. The higher, the smoother the image, but also the longer the processing time. For preview, the default of 16 samples should be sufficient and is also the fastest.

No zbuffer
: Sometimes you might want to have more control to blur the image. For instance, you may want to only blur one object while leaving everything else alone (or the other way around), or you want to blur the whole image uniformly all at once. The node therefore allows you to use something other than an actual zbuffer as the Z input. For instance, you could connect an image node and use a grayscale image where the color designates how much to blur the image at that point, where white is maximum blur and black is no blur. Or, you could use a Time node to uniformly blur the image, where the time value controls the maximum blur for that frame. It may also be used to obtain a possibly slightly-better DoF blur, by using a fake depth shaded image instead of a zbuffer. (A typical method to create the fake depth shaded image is by using a linear blend texture for all objects in the scene or by using the 'fog/mist' fake depth shading method.) This also has the advantage that the fake depth image can have anti-aliasing, which is not possible with a real zbuffer. "No zbuffer" will be enabled automatically whenever you connect a node that is not image based (e.g. time node/value node/etc).

Zscale
: Only visible when No zbuffer enabled. When No zbuffer is used, the input is used directly to control the blur radius. And since usually the value of a texture is only in the numeric range 0.0 to 1.0, its range is too narrow to control the blur properly. This parameter can be used to expand the range of the input (or for that matter, narrow it as well, by setting it to a value less than one). So for No zbuffer, this parameter therefore then becomes the main blur control (similar to fStop when you *do* use a zbuffer).

Examples



In this blend file example, the ball array image is blurred as if it was taken by a camera with a f-stop of 2.8 resulting in a farily narrow depth of field centered on 7.5 blender units from the camera. As the balls receed into the distance, they get blurrier.

Hints

Preview

In general, use preview mode, change parameters to your liking, only then disable preview mode for the final render. This node is compute intensive, so watch your console window, and it will give you status as it computes each render scan line.

Edge Artifacts

For minimum artifacts, try to setup your scene such that differences in distances between two objects that may visibly overlap at some point are not too large.

"Focus Pull"

Keep in mind that this is not 'real' DoF, only a post-processing simulation. Some things cannot be done which would be no problem for real DoF at all. A typical example is a scene with some object very close to the camera, and the camera focusing on some point far behind it. In the real world, using shallow depth of field, it is not impossible for nearby objects to become completely invisible, in effect allowing the camera to see 'behind' it. Hollywood cinematographers use this visual characteristic to good effect to achieve the popular "focus pull" effect, where the focus shifts from a nearby to a distant object, such that the "other" object all but disappears. Well, this is simply not possible to do with the current post-processing method in a single pass. If you really want to achieve this effect, quite satisfactorily, here's how:

- ◊ Split up your scene into "nearby" and "far" objects, and render them in two passes.
- ◊ Now, combine the two the two results, each with their own "defocus" nodes driven by the same Time node, but with one of them inverted. (e.g. using a "Map Value" node with a Size of -1.) As the defocus of one increases, the defocus on the other decreases at the same rate, creating a smooth transition.

Aliasing at Low f-Stop Values

At very low values, less than 5, the node will start to remove any oversampling and bring the objects at DoFDist very sharply into focus. If the object is against a constrasting background, this may lead to visible stairstepping (aliasing) which OSA is designed to avoid. If you run into this problem:

- ◊ Do your own OSA by rendering at twice the intended size and then scaling down, so that adjacent pixels are blurred togther
- ◊ Use the blur node with a setting of 2 for x and y
- ◊ Set DoFDist off by a little, so that the object in focus is blurred by the tiniest bit.
- ◊ Use a higher f-Stop, which will start the blur, and then use the Z socket to a Map Value to a Blur node to enhance the blur effect.
- ◊ Rearrange the objects in your scene to use a lower-contrast background

No ZBuffer

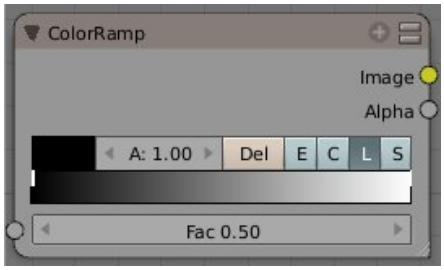
A final word of warning, since there is no way to detect if an actual zbuffer is connected to the node, be VERY careful with the "No ZBuffer" switch. If the Zscale value happens to be large, and you forget to set it back to some low value, the values may suddenly be interpreted as huge blur-radius values that will cause processing times to explode.

As the name implies, these nodes convert the colors or other propertiers of various data (e.g. images) in some way. They also split out or re-combine the different color channels that make up an image, allowing you to work on each channel independently. Various color channel arrangements are supported, including traditional RGB and HSV formats, and the newest High Definition Media Interface (HDMI) formats.

ColorRamp Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â ColorRamp



ColorRamp node
The ColorRamp Node is used for mapping values to colors with the use of a gradient. It works exactly the same way as a colorband for textures and materials, using the Factor value as a slider or index to the color ramp shown, and outputting a color value and an alpha value from the output sockets.

By default, the ColorRamp is added to the node map with two colors at opposite ends of the spectrum. A completely black black is on the left (Black as shown in the swatch with an Alpha value of 1.00) and a whitewash white is on the right. To select a color, LMB click on the thin vertical line/band within the colorband. The example picture shows the black color selected, as it is highlighted white. The settings for the color are shown above the colorband as (left to right): color swatch, Alpha setting, and interpolation type.

To change the hue of the selected color in the colorband, LMB click on the swatch, and use the popup color picker control to select a new color. Press â µ Enter to set that color.

To add colors, hold Ctrl down and CtrlLMB click inside the gradient. Edit colors by clicking on the rectangular color swatch, which pops up a color-editing dialog. Drag the gray slider to edit A:lpha values. Note that you can use textures for masks (or to simulate the old "Emit" functionality) by connecting the alpha output to the factor input of an RGB mixer.

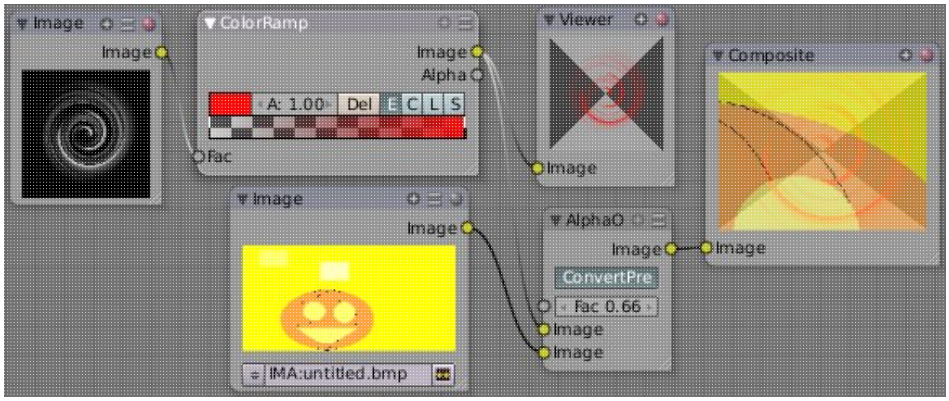
To delete a color from the colorband, select it and press the Delete button.

When using multiple colors, you can control how they transition from one to another through an interpolation mixer. Use the interpoloation buttons to control how the colors should band together: Ease, Cardinal, Linear, or Spline.

Use the A: button to define the Alpha value of the selected color for each color in the range.

Using ColorRamp to create an Alpha Mask

A powerful but often overlooked feature of the ColorRamp is to create an Alpha Mask, or a mask that is overlaid on top of another image, and, like a mask, allows some of the background to show through. The example map below shows how to use the Color Ramp node to do this:



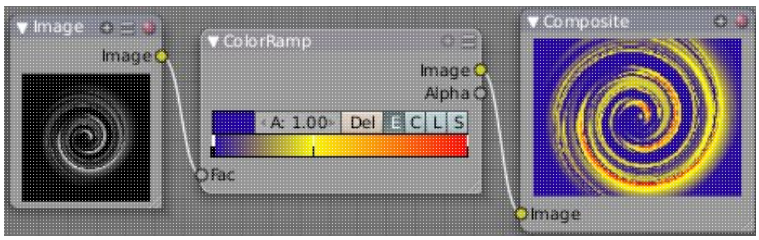
Using the ColorRamp node to create an alpha mask
In the map above, a black and white swirl image, which is lacking an alpha channel, is fed into the ColorRamp node as a Factor. (Technically, we should have converted the image to a value using the RGB-to-BW node, buy hey, this works just as well since we are using a BW image as input.)

We have set the ColorRamp node to a purely transparent color on the left end of the spectrum, and a fully Red color on the right. As seen in the viewer, the ColorRamp node puts out a mask that is fully transparent where the image is black. Black is zero, so ColorRamp uses the 'color' at the left end of the spectrum, which we have set to transparent. The ColorRamp image is fully red and opaque where the image is white (1.00).

We verify that the output image mask is indeed transparent by overlaying it on top of a pumpkin image. For fun, we made that AlphaOver output image 0.66 transparent so that we can, in the future, overlay the image on a flashing white background to simulate a scary scene with lighting flashes.

Using ColorRamp to Colorize an Image

The real power of ColorRamp is that multiple colors can be added to the color spectrum. This example compositing map takes a boring BW image and makes it a flaming swirl!

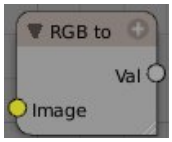


In this example, we have mapped the shades of gray in the input image to three colors, blue, yellow, and red, all fully opaque (Alpha of 1.00). Where the image is black, ColorRamp substitutes blue, the currently selected color. Where it is some shade of gray, ColorRamp chooses a corresponding color from the spectrum (bluish, yellow, to redish). Where the image is fully white, ColorRamp chooses red.

RGB to BW Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â RGB to BW

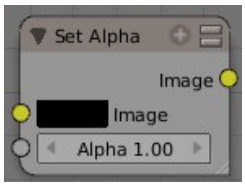


RGB to BW node
This node converts a color image to black-and-white.

Set Alpha Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â Set Alpha



Set Alpha node
This node adds an alpha channel to a picture. Some image formats, such as JPEG, do not support an alpha channel. In order to overlay a JPEG image on top of a background, you must add an alpha channel to it using this node.

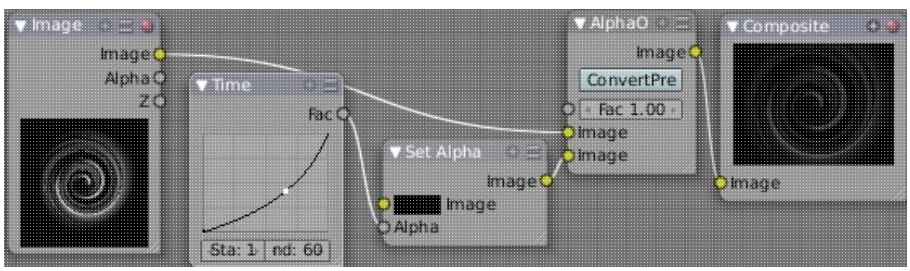
The Image input socket is optional. If an input image is not supplied, the base color shown in the swatch will be used. To change the color, LMB click the swatch and use the color-picker control to choose or specify a color you want.

The amount of Alpha (1.00 being totally opaque and 0.00 being totally transparent) can be set for the whole picture using the input field. Additionally, the Alpha factor can be set by feeding its socket.

Note that this is not, and is not intended to be, a general-purpose solution to the problem of compositing an image that doesn't contain Alpha information. You might wish to use "Chroma Keying" or "Difference Keying" (as discussed elsewhere) if you can. This node is most often used (with a suitable input being provided by means of the socket) in those troublesome cases when you *can't*, for some reason, use those techniques directly.

Using SetAlpha to Fade to Black

To transition the audience from one scene or shot to another, a common technique is to "fade to black". As its name implies, the scene fades to a black screen. You can also "fade to white' or whatever color you wish, but black is a good neutral color that is easy on the eyes and intellectually "resets" the viewer's mind. The node map below shows how to do this using the Set Alpha node.

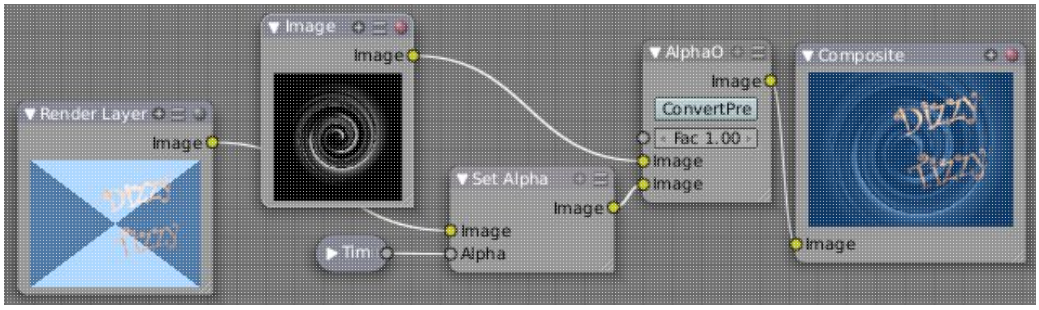


Fade To Black
In the example above, the alpha channel of the swirl image is ignored. Instead, a time node introduces a factor from 0.00 to 1.00 over 60 frames, or about 2 seconds, to the Set Alpha node. Note that the time curve is exponentially-shaped, so that the overall blackness will fade in slowly and then accelerate toward the end. The Set Alpha node does not need an input image; instead the flat (shadeless) black color is used. The Set Alpha Node uses the input factor and color to create a black image that has an alpha set which goes from 0.00 to 1.00 over 60 frames, or completely transparent to completely opaque. Think of alpha as a multiplier for how vivid you can see that pixel. These two images are combined by our trusty AlphaOver node completely (a Factor of 1.00) to produce the composite image. The SetAlpha node will thus, depending on the frame being rendered, produce a black image that has some degree of transparency. Set up and Animate, and you have an image sequence that fades to black over a 2-second period.

No Scene information used
This example node map does not use the RenderLayer. To produce this 2 second animation, no blender scene information was used. This is an example of using Blender's powerful compositing abilities separate from its modeling and animation capabilities. (A Render Layer could be substituted for the Image layer, and the "fade-network" effect will still produce the same effect)

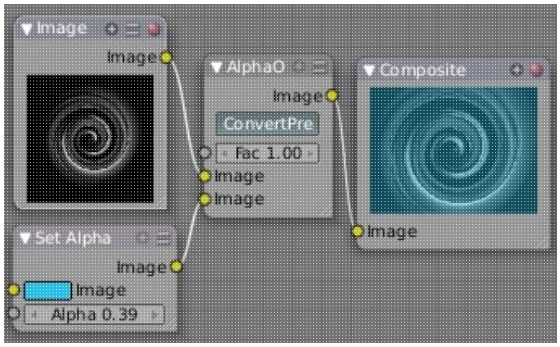
Using SetAlpha to Fade In a Title

To introduce your animation, you will want to present the title of your animation over a background. You can have the title fly in, or fade it in. To fade it in, use the SetAlpha node with the Time node as shown below.



Using Set Alpha to Fade in a Title
In the above example, a Time curve provides the Alpha value to the input socket. The current RenderLayer, which has the title in view, provides the image. As before, the trusty AlphaOver node mixes (using the alpha values) the background swirl and the alphaed title to produce the composite image. Notice the ConvertPre-Multiply button is NOT enabled; this produces a composite where the title lets the background image show through where even the background image is transparent, allowing you to layer images on top of one another.

Using SetAlpha to Colorize a BW Image



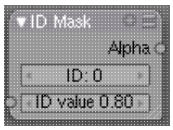
Using Set Alpha to Colorize an Image
In the example above, notice how the blue tinge of the render input colors the swirl. You can use the Set Alpha node's color swatch with this kind of node map to add a consistent color to a BW image.

In the example map to the right, use the Alpha value of the SetAlpha node to give a desired degree of colorization. Thread the input image and the Set Alpha node into an AlphaOver node to colorize any black and white image in this manner. Note the ConvertPre-Multiply button is enabled, which tells the AlphaOver node not to multiply the alpha values of the two images together.

ID Mask Node

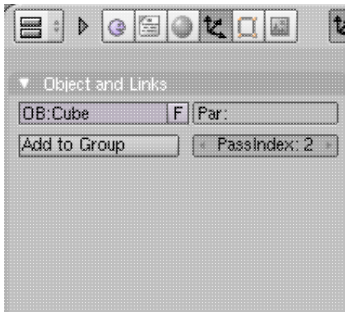
Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â ID Mask



ID Mask node
This node will use the Object Index pass (see RenderLayers) to produce an anti-aliased alpha mask for the object index specified. The mask is opaque where the object is, and transparent where the object isn't. If the object is partially transparent, the alpha mask matches the object's transparency. This post-process function fills in the jaggies with interpolated values.

Object Index
Object indices are only output from a RenderLayers node or stored in a multilayer OpenEXR format image.

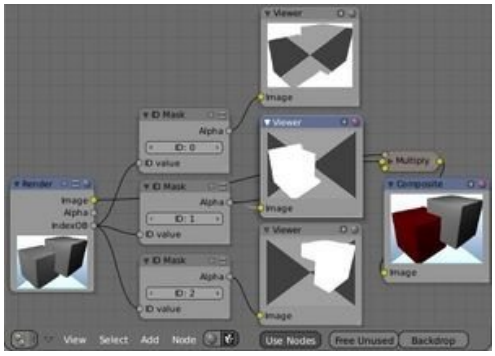


Setting an Object Index

You can specify, for any of the objects in your scene, an Object Index as shown the right (the currently select object has an index of 2). When rendered, if Object Index passes are enabled, its index will be 2, and setting the ID Mask node to 2 will show where that object is in the scene.

This node is extreeeemely well suited to removing the aliases shown as output from the Defocus node or DOF noodles caused by some objects being close to camera against objects far away.

Example



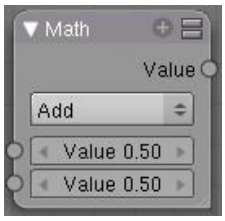
Example
In this example, the left rear red cube is assigned PassIndex 1, and the right cube PassIndex 2. Where the two cubes intersect, there is going to be noticeable pixelation (jaggies) because they come together at a sharp angle and are different colors. Using the mask from object 1, which is smoothed (anti-aliased) at the edges, we use a Mix node set on Multiply to multiply the smoothed edges against the image, thus removing those nasty (Mick) Jaggies. Thus, being smoothed out, the Rolling Stones gather no moss. (I really hope you get that obscure reference :)

Note that the mask returns white where the object is fully visible to the camera(not behind anything else) and black for the part of the object that is partially or totally obscured by a fully or partially opaque object in front of it. If something else is in front of it, even if that thing is partially transparent and you can see the object in a render, the mask will not reflect that partially obscured part.

Math Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â Math



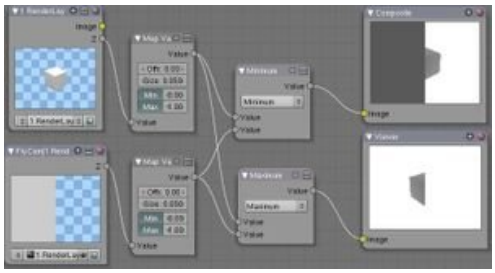
Math node
This node performs the selected math operation on an image or buffer. All common math functions are supported. If only an image is fed to one Value socket, the math function will apply the other Value consistently to every pixel in producing the output Value. Select the math function by clicking the up-down selector where the "Add" selection is shown.

The trig functions of Sine, Cosine, Tangent use only the top socket and accept values in radians between 0 and 2*pi for one complete cycle.

Blender 2.44+
Known bug: the Top socket must get the image if the bottom socket is left as a value.

Examples

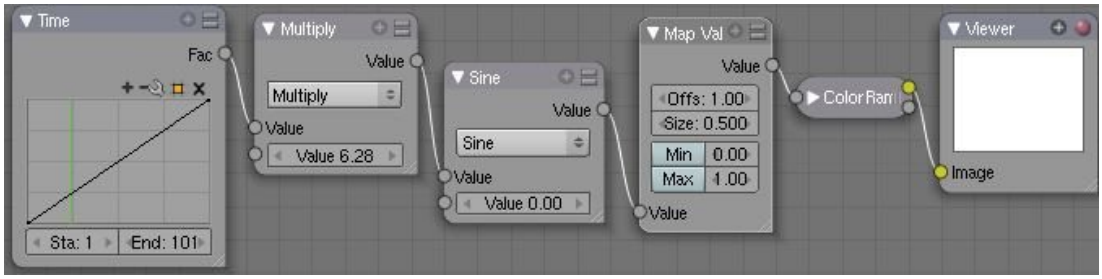
Manual Z-Mask



Example
This example has one scene input by the top RenderLayer node, which has a cube that is about 10 BU from the camera. The bottom RenderLayer node inputs a scene (FlyCam) with a plane that covers the left half of the view and is 7 BU from the camera. Both are fed through their respective Map Value nodes to divide the Z buffer by 20 (multiply by .05, as shown in the Size field) and clamped to be a Min/Max of 0.0/1.0 respectively.

For the Minimum function, the node selects those Z values where the corresponding pixel is closer to the camera; so it chooses the Z values for the plane and part of the cube. The background has an infinite Z value, so it is clamped to 1.0 (shown as white). In the maximum example, the Z values of the cube are greater than the plane, so they are chosen for the left side, but the plane (FlyCam) Renderlayer's Z are infinite (mapped to 1.0) for the right side, so they are chosen.

Using Sine Function to Pulsate



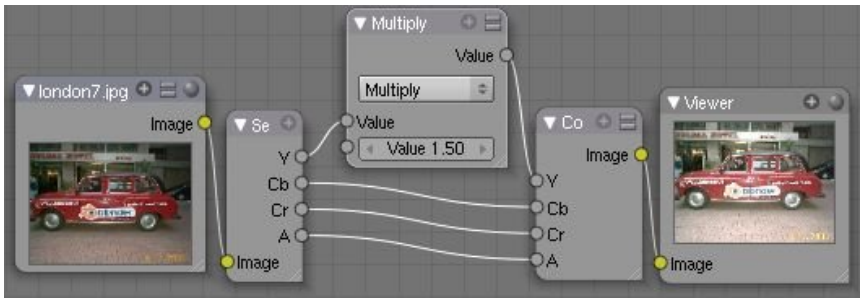
This example has a Time node putting out a linear sequence from 0 to 1 over the course of 101 frames. The green vertical line in the curve widget shows that frame 25 is being put out, or a value of .25. That value is multiplied by 2*pi and converted to 1.0 by the Sine function, since we all know that Sine(2*pi/4)=Sine(pi/2)=+1.0.

Since the Sine function can put out values between -1.0 and 1.0, the Map Value node scales that to 0.0 to 1.0 by taking the input (-1 to 1), adding 1 (making 0 to 2), and multiplying the result by one half (thus scaling the output between 0 and 1). The default ColorRamp converts those values to a grayscale. Thus, medium gray corresponds to a 0.0 output by the sine, black to -1.0, and white to 1.0. As you can see, Sine(pi/2)=1.0. Like having your own visual color calculator! Animating this noodle provides a smooth cyclic sequence through

the range of grays.

Use this function to vary, for example, the alpha channel of an image to produce a fading in/out effect. Alter the Z channel to move an scene in/out of focus. Alter a color channel value to make a color "pulse".

Brightening/Scaling a Channel



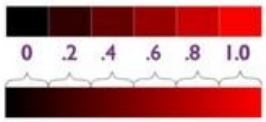
This example has a Multiply node increasing the luminance channel (Y) of the image to make it brighter. Note that you should use a Map Value node with Min() and Max() enabled to clamp the output to valid values. With this approach you could use a logarithmic function to make a high-dynamic range image. For this particular example, there is also a Brighten/Contrast node that might give simpler control over brightness.

Quantize/Restrict Color Selection

In this example, we want to restrict the color output to only 256 possible values. Possible use of this is to see what the image will look like on an 8-bit cell phone display. To do this, we want to restrict the R, G and B values of any pixel to be one of a certain value, such that when they are combined, will not result in more than 256 possible values. The number of possible values of an output is the number of channel values multiplied by each other, or Q = R * G * B.

Since there are 3 channels and 256 values, we have some flexibility how to quantize each channel, since there are a lot of combinations of R*G*B that would equal 256. For example, if {R,G,B} = {4,4,16}, then 4 * 4 * 16 = 256. Also, {6,6,7} would give 252 possible values. The difference in appearance between {4,4,16} and {6,6,7} is that the first set (4,4,16} would have fewer shades of red and green, but lots of shades of blue. The set {6,6,7} would have a more even distribution of colors. To get better image quality with fewer color values, give more possible values to the predominant colors in the image.

Theory



Two Approaches to Quantizing to 6 values
To accomplish this quantization of an image to 256 possible values, lets use the set {6,6,7}. To split up a continuous range of values between 0 and 1 (the full Red spectrum) into 6 values, we need to construct an algorithm or function that takes any input value but only puts out 6 possible values, as illustrated by the image to the right. We want to include 0 as true black, with five other colors in between. The approach shown produces {0,.2,.4,.6,.8,1}. Dividing 1.0 by 5 equals .2, which tells us how far apart each quantified value is from the other.

So, to get good even shading, we want to take values that are 0.16 or less and map them to 0.0; values between 0.16 and 0.33 get fixed to 0.2; colorband values between 0.33 and 0.5 get quantized to 0.4, and so on up to values between 0.83 and 1.0 get mapped to 1.0.

Function f(x)
An algebraic function is made up of primitive mathematical operations (add, subtract, multiply, sine, cosine, etc) that operate on an input value to provide a desired output value.

Input		Function				Output	
8-bit	R value	*n	- ½	round()	/(n-1)	R value	8-bit
0	0.00	0.00	-0.50	0	0	0	0
13	0.05	0.30	-0.20	0	0	0	0
26	0.10	0.60	0.10	0	0	0	0
38	0.15	0.90	0.40	0	0	0	0
51	0.20	1.20	0.70	1	0.2	0.2	51
64	0.25	1.50	1.00	1	0.2	0.2	51
77	0.30	1.80	1.30	1	0.2	0.2	51
89	0.35	2.10	1.60	2	0.4	0.4	102
102	0.40	2.40	1.90	2	0.4	0.4	102
115	0.45	2.70	2.20	2	0.4	0.4	102
128	0.50	3.00	2.50	2	0.4	0.4	102
140	0.55	3.30	2.80	3	0.6	0.6	153
153	0.60	3.60	3.10	3	0.6	0.6	153
166	0.65	3.90	3.40	3	0.6	0.6	153
179	0.70	4.20	3.70	4	0.8	0.8	204
191	0.75	4.50	4.00	4	0.8	0.8	204
204	0.80	4.80	4.30	4	0.8	0.8	204
217	0.85	5.10	4.60	5	1	1	255
230	0.90	5.40	4.90	5	1	1	255
242	0.95	5.70	5.20	5	1	1	255
255	1.00	6.00	5.50	5	1	1	255

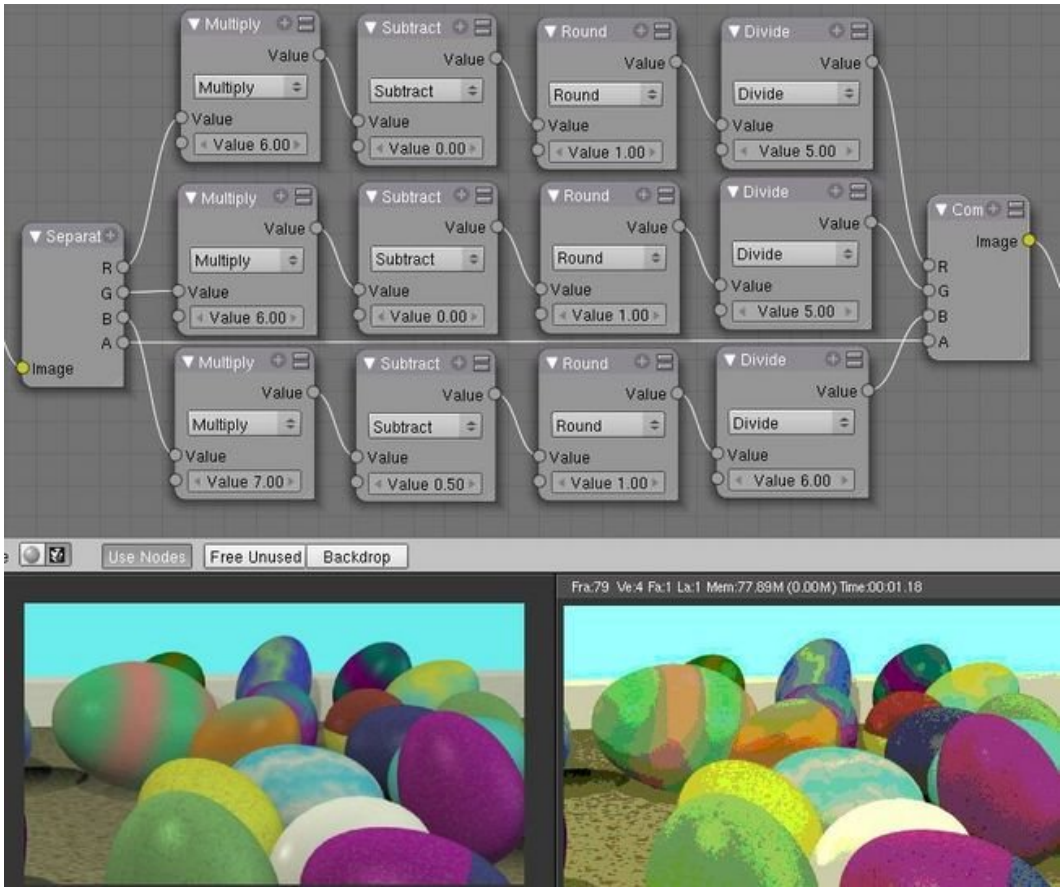
The theory behind this function is scaled truncation. Let us suppose we want a math function that takes in a range of values between 0 and 1, such as .552, but only outputs a value of 0.0, 0.2, 0.4, etc. We can imagine then that we need to get that range 0 to 1 powered up to something 0 to 6 so that we can chop off and make it a whole number. So, with six divisions, how can we do that? The answer is we multiply the range by 6. The output of that first math multiply node is a range of values between 0 and 6. To get even divisions, because we are using the rounding function (see documenation above), we want any number plus or minus around a whole number will get rounded to that number. So, we subtract a half, which shifts everything over. The Round() function then makes that range 0 to 5. We then divide by 5 to get back a range of numbers between 0 and 1 which can then be combined back with the other color channels. Thus, you get the the function

f(x,n)=round[x*n-1/2] /(n-1)

where n is the number of possible output values, and x is the input pixel color and f(x,n) is the output value. There's only one slight problem, and that is for the value exactly equal to 1, the formula result is 1.2, which is an invalid value. This is because the round function is actually a roundup function, and exactly 5.5 is rounded up to 6. So, by subtracting .501, we compensate and thus 5.499 is rounded to 5. At the other end of the spectrum, pure black, or 0, when .501 subtracted, rounds up to 0 since the Round() function does not return a negative number.

Sometimes using a spreadsheet can help you figure out how to put these nodes together to get the result that you want. Stepping you through the formula for n=6 and x=0.70, locate the line on the spreadsheet that has the 8-bit value 179 and R value 0.7. Multiplying by 6 gives 4.2. Subtracting 1/2 gives 3.7, which rounds up to 4. 4 divided by 5 = .8. Thus, f(0.7, 6) = 0.8 or an 8-bit value of 204. You can see that this same 8-bit value is output for a range of input values. Yeah! Geeks Rule! This is how you program Blender to do compositing based on Algebra. Thank a Teacher if you understand this.

Reality



To implement this function in Blender, consider the noodle above. First, feed the image to the Separate RGB node. For the Red channel, we string the math nodes into a function that takes each red color, multiplies (scales) it up by the desired number of divisions (6), offsets it by 0.5, rounds the value to the nearest whole number, and then divides the image pixel color by 5. So, the transformation is {0..1} becomes {0..6}, subtracting centers the medians to {-0.5...5.5} and the rounding to the nearest whole number produces {0,1,2,3,4,5} since the function rounds down, and then dividing by five results in six values {0.0,0.2,0.4,0.6,0.8,1.0}.

The result is that the output value can only be one of a certain set of values, stair-stepped because of the rounding function of the math node noodle. Copying this one channel to operate on Green and Blue gives the noodle below. To get the 6:6:7, we set the three multiply nodes to {6,6,7} and the divide nodes to {5,5,6}.

If you make this into a node group, you can easily re-use this setup from project to project. When you do, consider using a math node to drive the different values that you would have to otherwise set manually, just to error-proof your work.

Summary

Normally, an output render consists of 32- or 24-bit color depth, and each pixel can be one of millions of possible colors. This noodle example takes each of the Red, Green and Blue channels and normalizes them to one of a few values. When all three channels are combined back together, each color can only be one of 256 possible values.

While this example uses the Separate/Combine RGB to create distinct colors, other Separate/Combine nodes can be used as well. If using the YUV values, remember that U and V vary between -0.5 and +0.5, so you will

have to first add on a half to bring the range between 0 and 1, and then after dividing, subtract a half to bring in back into standard range.

The JPG or PNG image format will store each of the colors according to their image standard for color depth (e.g. JPG is 24-bit), but the image will be very very small, since reducing color depth and quantizing colors is essentially what the JPEG compression algorithm accomplishes.

You do not have to reduce the color depth of each channel evenly. For example, if blue was the dominant color in an image, to preserve image quality, you could reduce Red to 2 values, Green to 4, and let the blue take on 256/(2*4) or 32 values. If using the HSV, you could reduce the Saturation and Value to 2 values (0 or 1.0) by Multiply by 2 and Divide by 2, and restrict the Hue to 64 possible values.

You can use this noodle to quantize any channel; alpha, speed (vector), z-values, and so forth.

Combine/Separate Nodes

All of these node do essentially the same thing: they split out an image into (or recombine an image from) its composite color channels. Each format supports the Alpha (transparency) channel. The standard way of representing color in an image is called a *color space*. There are several color spaces supported:

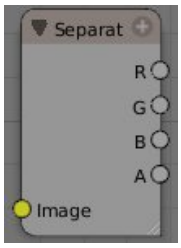
- RGB: Red-Green-Blue traditional primary colors, also broadcast directly to most computer monitors
- HSV: Three values, often considered as more intuitive than the RGB system (nearly only used on computers):
 - Hue: the **Hue** of the color (in some way, choose a 'color' of the rainbow);
 - Saturation: the **quantity** of hue in the color (from desaturate - shade of gray - to saturate - brighter colors);
 - Value: the **luminosity** of the color (from 'no light' - black - to 'full light' - 'full' color, or white if Saturation is 0.0).
- YUV: Luminance-Chrominance standard used in broadcasting analog PAL (European) video.
- YCbCr: Luminance-ChannelBlue-ChannelRed Component video for digital broadcast use, whose standards have been updated for HDTV and commonly referred to as the HDMI format for component video.

See the global wikipedia for more information on color spaces.

Separate/Combine RGBA Node

Panel: Node Editor → Node Composition

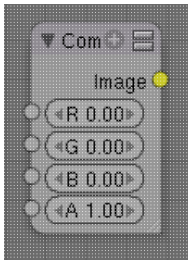
Menu: â § ShiftA â Convertors â Separate RGBA



Separate RGBA node
This node separates an image into its red, green, blue and alpha channels. There's a socket for each channel on the right.

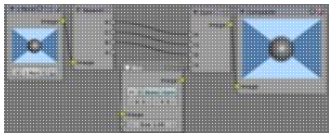
Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â Combine RGBA

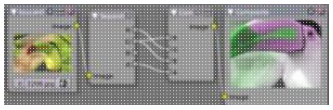


Combine RGBAnode
This node combines separate input images as each color and alpha channel, producing a composite image. You use this node combine the channels after working on each color channel separately.

Examples



In this first example, we take the Alpha channel and blur it, and then combine it back with the colors. When placed in a scene, the edges of it will blend in, instead of having a hard edge. This is almost like anti-aliasing, but in a three-dimensional sense. Use this noodle when adding CG elements to live action to remove any hard edges. Animating this effect over a broader scale will make the object appear to "phase" in and out, as a "out-of-phase" time-traveling sync effect.

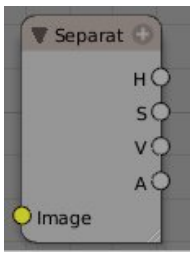


In this fun little noodle we make all the reds become green, and all the green both Red and Blue, and remove Blue from the image completely. Very cute. Very fun.

Separate/Combine HSVA Nodes

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â Separate HSVA



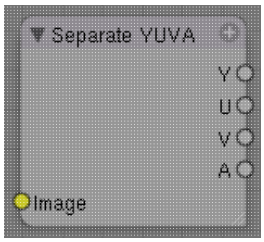
Separate HSVA node
This node separates an image into image maps for the hue, saturation, value and alpha channels.

Use and manipulate the separated channels for different purposes; i.e. to achieve some compositing/color adjustment result. For example, you could expand the Value channel (by using the multiply node) to make all the colors brighter. You could make an image more relaxed by diminishing (via the divide or map value node) the Saturation channel. You could isolate a specific range of colors (by clipping the Hue channel via the Colorramp node) and change their color (by the Add/Subtract mix node).

Separate/Combine YUVA Node

Panel: Node Editor → Node Composition

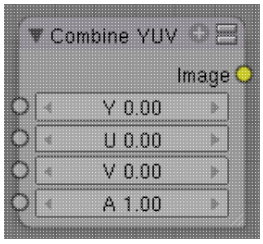
Menu: â § ShiftA â Convertors â Separate YUVA



Separate YUVA node
This node converts an RGBA image to YUVA color space, then splits each channel out to its own output so that they can be manipulated independently. Note that U and V values range from -0.5 to +0.5.

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â Combine YUVA

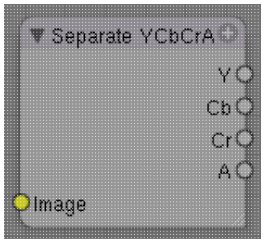


Combine YUVA node
Combines the channels back into a composite image. If you do not connect any input socket, you can set a default value for the whole image for that channel using the numeric controls shown.

Separate/Combine YCbCrA Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â Separate YCbCrA



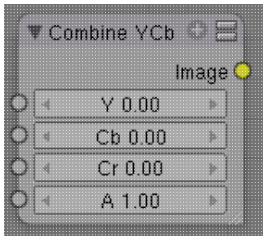
Separate YCbCrA node
This node converts an RGBA image to YCbCrA color space, then splits each channel out to its own output so that they can be manipulated independently:

- Y: Luminance, 0=black, 1=white
- Cb: Chrominance Blue, 0=Blue, 1=Yellow
- Cr: Chrominance Red, 0=Red, 1=Yellow

Note: If running these channels through a ColorRamp to adjust value, use the Cardinal scale for accurate representation. Using the Exponential scale on the luminance channel gives high-contrast effect.

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Convertors â Combine YCbCrA



Combine YCbCrA node
So, I kinda think you get the idea, and I was trying to think of some other creative way to write down the same thing, but I can't. So, you'll have to figure this node out on your own.

These nodes give you the essential tools for working with blue-screen or green-screen footage, where live action is shot in front of a blue or green backdrop for replacement by a matte painting or virtual background.

In general, hook up these nodes to a viewer, set your UV/Image Editor to show the viewer node, and play with the sliders in real-time using a sample image from the footage, to get the settings right. In some cases, small adjustments can eliminate artifacts or foreground image degredation. For example, taking out too much green can result in foreground actors looking 'flat' or blueish/purplish.

You can and should chain these nodes together, refining your color correction in successive refinements, using each node's strengths to operate on the previous node's output. There is no "one stop shopping" or one "does-it-all" node; they work best in combination.

Usually, green screen is shot in a stage with consistent lighting from shot to shot, so the same settings will work across multiple shots of raw footage. Footage shot outside under varying lighting conditions (and wind blowing the background) will complicate matters and mandate lower falloff values.

Garbage Matte

Garbage matte is not a node, but a technique where the foreground is outlined using a closed curve (bezier or nurbs). Only the area within the curve is processed using these matte nodes; everything else is garbage and thus discarded.

Difference Key Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Mattes â Difference Key



Difference Key node
The difference key node determines how different each channel is from the given key in the selected color space. If the differences are below a user defined threshold then the pixel is considered transparent. Difference matting does not rely on a certain background color, but can have less than optimal results if there is a significant amount of background color in the foreground object.

There are two inputs to this node.

- The first is an input Image that is to be keyed.
- The Key Color can be input as an RGB value or selected using the color picker by clicking on the Key Color box to bring up the color dialog, then clicking on the eye dropper tool and selecting a color.

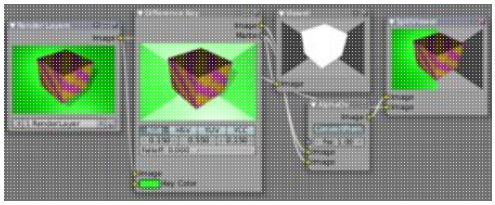
The selectable color spaces are RGB (default), HSV, YUV, and YCbCr.

You can adjust the tolerance of each color in the colorspace individually so that you can have more red variance or blue variance in what you would allow to be transparent. I find that about 0.15 (or 15%) is plenty of variance if the background is evenly lit. Any more unevenness and you risk cutting into the foreground image.

When the Falloff value is high, pixels that are close to the Key Color are more transparent than pixels that are not as close to the Key Color (but still considered close enough to be keyed). When the Falloff value is low, it does not matter how close the pixel color (Image) is to the Key Color, it is transparent.

The outputs of this node are the Image with an alpha channel adjusted for the keyed selection and a black and white Matte (i.e the alpha mask).

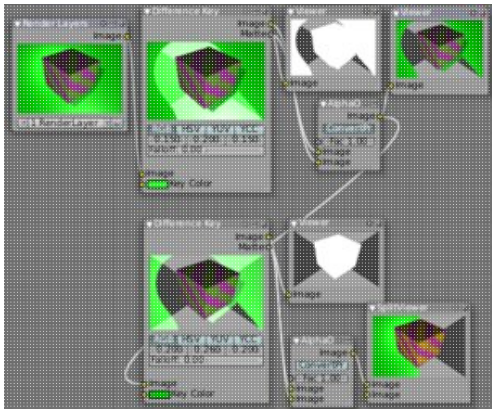
Simple Example



Using the Difference Key Node
In the example to the right (click to expand), we have a purple cube with yellow marbeling in front of a very unevenly lit green screen. We start building our noodle by threading the image to a difference key, and using the eyedropper, pick a key color very close to the edge of the cube, around where the halo is at the corner on the left-hand side; a fairly bright green. We thread two viewers from the output sockets so we can see what (if anything) the node is doing. We add an AlphaOver node, threading the Matte to the **TOP** socket and the image to the **BOTTOM** socket. Very Important, because 0 time blue is not the same as blue times zero. You always want your mask to go to the top socket of the AlphaOver. Premultiply is set and a full multiply is on so that we completely remove the green. In this example, we thread the output of the alphaover to a SplitViewer node so we can compare our results; the original is threaded to the bottom input of the SplitViewer, so that original is on the left, processed is on the right.

We set our variance to .15, and see what we get. What we get (not shown) is a matte that masks around the cube, but not on the right and around the edges where the green is darker; that shade it is too far away from our key color. So, since it is the green that is varying that we want to remove, we increase the Green variation to 1.00 (not shown). Whoa! All the Green disappears (all green within a 100% variation of our green key color is *all* the green), along with the top of the box! Not good. So, we start decreasing the green until we settle on 55% (shown).

Chaining Example



Chaining Difference Key Nodes
We pay out the wazoo for our highly talented (and egotistical I might add) Mr. Cube to come into the studio and do a few takes. We told him NOT to wear a green tie, but when we look at our footage, lo and behold, there he is with a green striped tie on. When we use our simple noodle, the green stripes on his tie go alpha, and the beach background shows through. So, we call him up and, too late, he's on his way back to Santa Monica and it wasn't in his contract and it wasn't his fault, after all, we're supposed to have all this fancy postpro software yada yada and he hangs up. Geez, these actors.

So, we chain two Difference Key nodes as shown to the right, and problem solved. What we did was lower the variation percentage on the first to remove some of the green, then threaded that to a second (lower)

difference key, where we sampled the green more toward the shadow side and outside edge. By keeping both variations low, none of the green in his tie is affected; that shade is outside the key's +/- variation tolerances.

Chroma Key Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Mattes â Chroma Key



Chroma Key node
The Chroma Key node determines if a pixel is foreground or background (and thereby should be transparent) based on its chroma values. This is useful for compositing images that have been shot in front of a green or blue screen.

There is one input to this node, the Image that is to be keyed.

Control this node using:

Green / Blue buttons
: Basic selection of what color the background is supposed to be.

Cb Slope and Cr Slope (chroma channel) sliders
: Determines how quickly the processed pixel values go from background to foreground, much like falloff.

Cb Pos and Cr Pos sliders
: Determines where the processing transition point for foreground and background is in the respective channel.

Threshold
: Determines if additional detail is added to the pixel if it is transparent. This is useful for pulling shadows from an image even if they are in the green screen area.

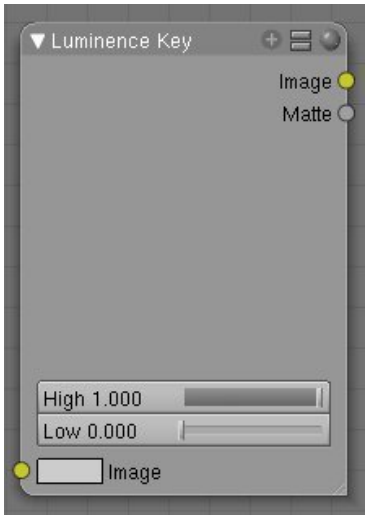
Alpha threshold
: The setting that determines the tolerance of pixels that should be considered transparent after they have been processed. A low value means that only pixels that are considered totally transparent will be transparent, a high value means that pixels that are mostly transparent will be considered transparent.

The outputs of this node are the Image with an alpha channel adjusted for the keyed selection and a black and white Matte (i.e the alpha mask).

Luminance Key Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Mattes â Luma Key



Luminance Key node
The Luminance Key node determines background objects from foreground objects by the difference in the luminance (brightness) levels. For example, this is useful when compositing stock footage of explosions (very bright) which are normally shot against a solid, dark background.

There is one input to this node, the Image that is to be keyed.

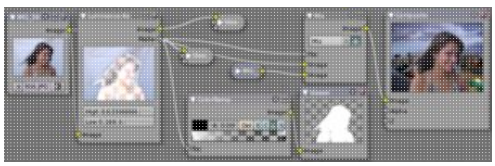
Control this node using:

- The High value selector determines the lowest values that are considered foreground. (which is supposed to be - relatively - light: from this value to 1.0).
- The Low value selector determines the hightes values that are considered to be background objects. (which is supposed to be - relatively - dark: from 0.0 to this value).

It is possible to have a separation between the two values to allow for a gradient of transparency between foreground and background objects.

The outputs of this node are the Image with an alpha channel adjusted for the keyed selection and a black and white Matte (i.e the alpha mask).

Example



Using Luma Key...with a twist
For this example, let's throw you a ringer. Here, the model was shot against a *white* background. Using the Luminance Key node, we get a matte out where the background is white, and the model is black; the opposite

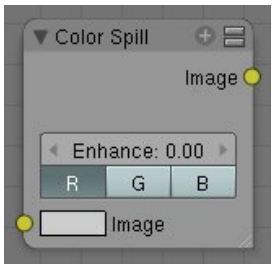
of what we want. If we wanted to use the matte, we have to switch the white and the black. How to do this? ColorRamp to the rescue - we set the left color White Alpha 1.0, and the right color to be Black Alpha 0.0. Thus, when the Colorramp gets in black, it spits out white, and vice versa. The reversed mask is shown; her white outline is useable as an alpha mask now.

Now to mix, we don't really need the AlphaOver node; we can just use the mask as our Factor input. In this kinda weird case, we can use the matte directly; we just switch the input nodes. As you can see, since the matte is white (1.0) where we don't want to use the model picture, we feed the background photo to the bottom socket (recall the mix node uses the top socket where the factor is 0.0, and the bottom socket where the factor is 1.0). Feeding our original photo into the top socket means it will be used where the Luminance Key node has spit out Black. Voila, our model is teleported from Atlanta to aboard a cruise ship docked in Miami.

Color Spill Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Mattes â Color Spill



Color Spill node
The Color Spill node reduces one of the RGB channels so that it is not greater than any of the others. This is common when compositing images that were shot in front of a green or blue screen. In some cases, if the foreground object is reflective, it will show the green or blue color; that color has "spilled" onto the foreground object. If there is light from the side or back, and the foreground actor is wearing white, it is possible to get "spill" green (or blue) light from the background onto the foreground objects, coloring them with a tinge of green or blue. To remove the green (or blue) light, you use this fancy node.

There is one input to this node, the Image to be processed.

The Enhance slider allows you to reduce the selected channel's input to the image greater than the color spill algorithm normally allows. This is useful for exceptionally high amounts of color spill.

The outputs of this node are the image with the corrected channels.

Channel Key Node

Blender 2.44+

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Mattes â Channel Key



Channel Key node
The Channel Key node determines background objects from foreground objects by the difference in the selected channel's levels. For example in YUV color space, this is useful when compositing stock footage of explosions (very bright) which are normally shot against a solid, dark background.

There is one input to this node, the Image that is to be keyed.

Control this node using:

- Color Space buttons selects what color space the channels will represent.
- Channel buttons selects the channel to use to determine the matte.
- High value selector determines the lowest values that are considered foreground. (which is supposed to be - relatively - hight values: from this value to 1.0).
- Low value selector determines the highest values that are considered to be background objects. (which is supposed to be - relatively - low values: from 0.0 to this value).

It is possible to have a seperation between the two values to allow for a gradient of transparency between foreground and background objects.

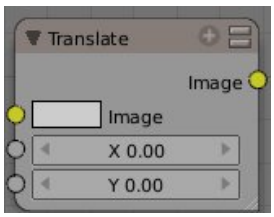
The outputs of this node are the Image with an alpha channel adjusted for the keyed selection and a black and white Matte (i.e the alpha mask).

These nodes distort the image in some fashion, operating either uniformly on the image, or by using a mask to vary the effect over the image.

Translate Node

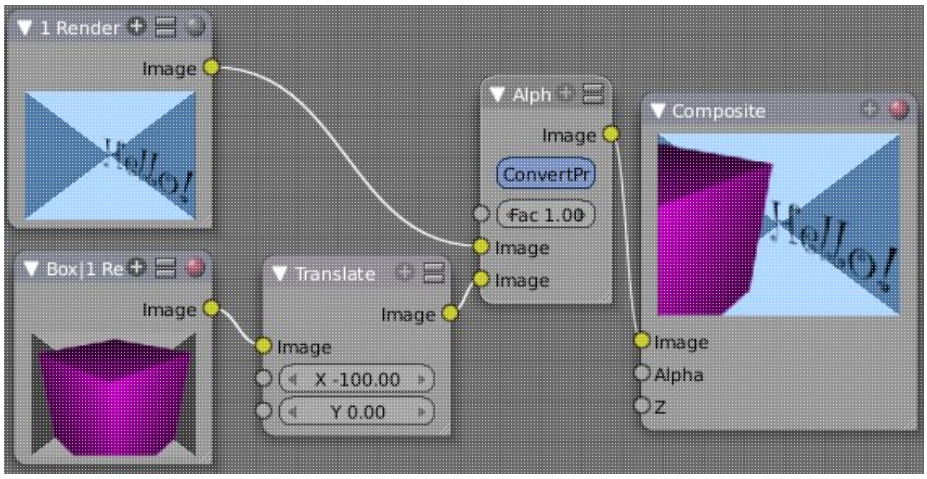
Panel: Node Editor → Node Composition

Menu: â § ShiftA â Distort â Translate



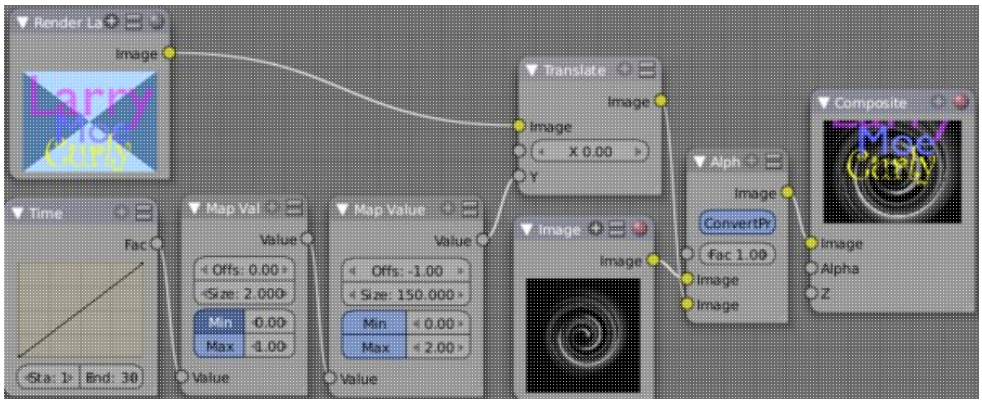
Translate node
The translate node translates (moves) an image by the specified amounts in the X and Y directions. X and Y are in pixels, and can be positive or negative. To shift an image up and to the left, for example, you would specify a negative X offset and a positive Y.

Use this node to position an image into the final composite image. Usually, the output of this node is routed to an AlphaOver or Mix node to mix it with a base image. In the example below, the RenderLayer input from one scene (box) is translated over to the left (a negative X translation) and alphaovered with a Hello scene RenderLayer input



Example: Using the Translate Node to Roll Credits

At the end of your fantastic animation, you'll want to give credit where credit is due. This is called rolling the credits and you see the names of everyone involved scroll up over a background image or sequence. The mini-map below shows an example of how to roll credits using the Translate node.



Using the Translate Node to Roll Credits

In this node map, the RenderLayer input has a list of the people involved and is 150 pixels high; it is the image input into the Translate Node. The Y value (vertical) offset of the Translate node comes from a scaled time factor that varies from -150 to 150 over 30 frames. The Translated image is overlaid on top of a background swirl image. So, over the course of 30 frames, the Translate node shifts the image from down by 150 pixels (off the bottom of the screen), up through and overlaid on top of the swirl, and finally off the screen to the top. These frames are generated when the Render Animation buttons are set and Anim is pressed. Right now, frame 21 is showing Moe and Curly, and Larry has scrolled off the screen.

Alpha channel

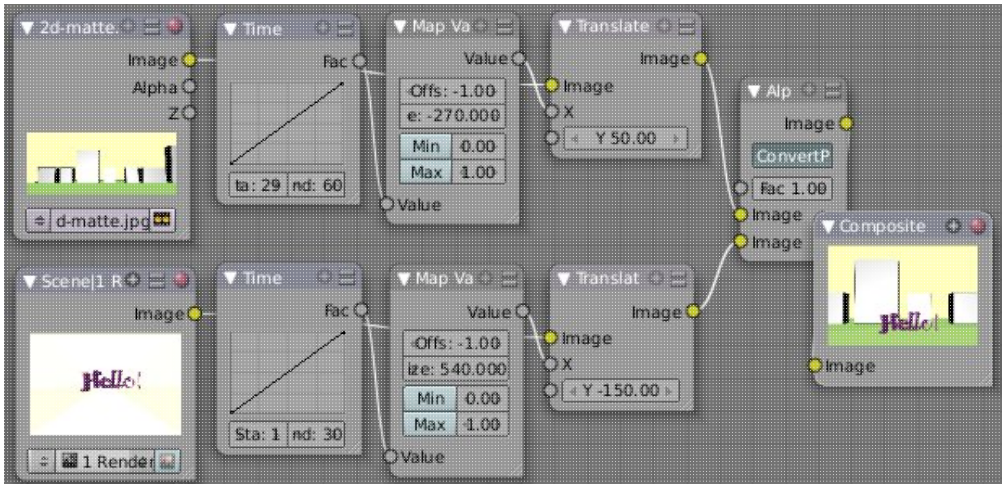
Be sure to save your credits image in a format that supports an alpha channel, such as PNG, so that the AlphaOver node can overlay it on the background and let the background show through.

You *could* have parented and animated all of the text to roll up past your camera, but rendering would have taken forever. Using the translate node is much faster to composite, and more flexible. To add someone to the list, simply change the credits image and re-animate. Since it is simple image manipulation, Blender is

blazingly fast at regenerating your credits. Similarly, changing the background is rock simple as well; simply load up a different background image and re-Animate.

Example: Moving a Matte

In some 2D and 3D animations and movies, a matte painting is used as the background. In most scenes it is still, however you can easily move it using the Translate node. Mattes are huge; the example used below is actually 1440x600 pixels, even though the scene being rendered is for TV. This oversizing gives us room to move the matte around. The example noodle below introduces a dancing "Hello!" from stage right in frames 1-30. As the "Hello" reaches center stage, we fake a camera move by moving the matte to the left, which makes it look like the "Hello!" is still moving as the camera moves with it.



Moving a Matte in back of a moving Animation (frame 60)
Use the Map Value node to scale the X (left to right) offsets that feed the Translate node. Note that offsets are used to position the dancing "Hello!" down to look like it is walking along the street (in the render scene, it is centered on camera and just dances in place). The matte is adjusted up to fake a camera height of an observer, bringing the horizon up.

Example: Shake, Rattle and Roll

A real world effect is the shaking of the camera. BOOM! We expect to feel the impact and see it rock our world. In our virtual CG world, though, we are in the vaccuum of space. So, we have to fake a camera being shook. There are two points in the development cycle to do this: at *Render-Scene* time, and at *Composite* time.

At *Render-Scene* time, the modeler would introduce an Ipo curve and keyframes that rotate and/or move the camera for a short amount of time. The advantage to render-scene shake is that the artist handles it and the editor does not have to worry about it; one less thing to do thank goodness. The disadvantage is that the artist may only be modeling the actors, and not the background scenery, props, or matte, so any shake they introduce does not transfer to the set, props, or backdrop. Therefore, it is best to introduce camera shake after all scenes have been rendered and assembled and when they are being composited.

There are two aspects to being bumped, or tripping over the tripod, or having an explosion go off next to you, or an airplane have a near miss as it flies by, or throwing up on a long sea voyage, or surviving an earthquake: *camera motion* and *image blur* (I know you were thinking expletives and changing your underpants, but this is about compositing).

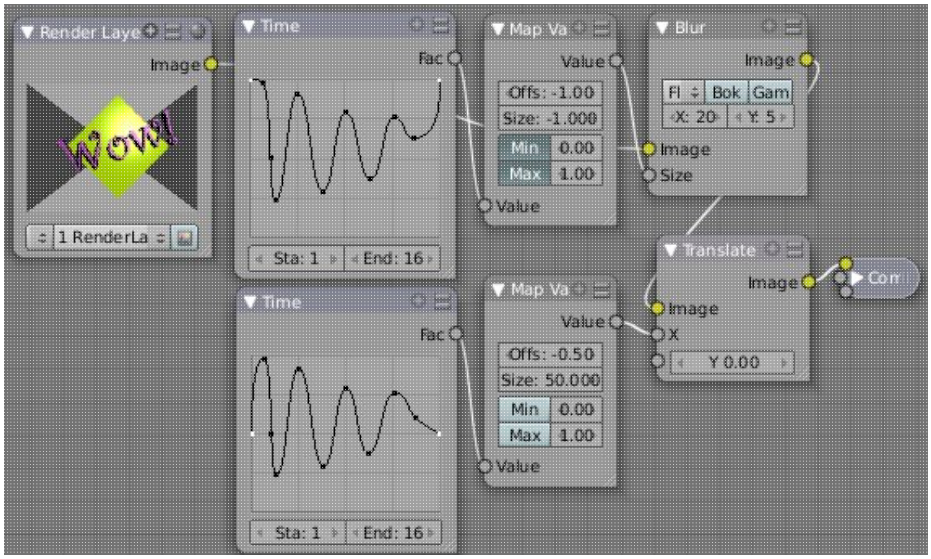
Camera Motion happens because the camera physically gets moved; but its mass and its tripod also acts as a

dampening device, softening out and absorbing the initial bump. The cameraman also acts as a dampener, and also as a corrector, trying to get the camera back to where it was pointed originally.

There can be quite a delay between the shock and the correction; for example, a lone actor/cameraman may trip on the tripod coming out from behind the camera, come into frame, realize the camera is off, and then come back to correct it. It all depends on the artistic effect and story you want to convey.

The *image blur* comes into play because the shake happens so rapidly that the image is blurred in the direction of the shake. However, the blur is more when the camera is being pushed back into position, and less when the camera is at the extreme of its deflection, since it is decellerating at the apex of its movement. Like motion, blur is the most during the initial shock, and less as things slow down and get under control. Also, the camera may go out of focus and come back into focus at the end of the shake.

To use Blender nodes to mimic Camera Motion, use the noodle shown below. The noodle has a Blur group on top that feeds a Translate group below it.



SFX: Camera Shake
In the above example, we use a Time curve that mimics the intensity and duration of a typical BOOM!. In this case, both curves have four peaks within a 16-frame period to mimic a BOOM! (in fact one curve was constructed and then duplicated to make the other, to ensure that the bulk of both curves was exactly the same). Notice how the curve dampens (decreases in magnitude as time progresses) as discussed above. Notice how the curve slows down (increasing period) to mimic the cameraman getting it back under control. Notice that the curve is sinusoidal to mimic over-correction and vibration.

BOOM! to the Left: The translate curve starts at 0.5. Maximum deflection up is fully a half, yet down is only a quarter. This offset mimics a BOOM! off to our left, since the camera shakes more to the right, away from the BOOM!

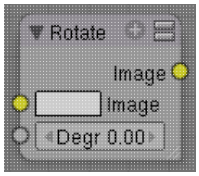
Motion and Blur are the same but different: Notice that the two curves are identical except for the highlighted start and end dots; we want zero blur and zero offsets before and after the shake, but minimum blur when there is maximum translate. The two Map Value node settings are different to achieve this; the math is left to the reader.

Use this Blender noodle to mimic camera shake. The amount of shake is set by the Size values, and the blur should be proportional to the amount and direction of motion (predominantly X in this example). Use the Time start and end to vary the duration of the shake; ten seconds for an earthquake, one minute for a ship rolling in the seas, a half second as an F-14 flies by and takes your ear off. *Author's note: I noticed cool camera shakes while watching the Halo 3 previews.*

Rotate Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Distort â Rotate

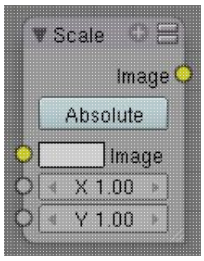


Rotate node
This node rotates an image. Positive values rotate clockwise and negative ones counterclockwise.

Scale Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Distort â Scale



Scale node
This node scales the size of an image. Scaling can be either absolute or relative. If Absolute toggle is on, you can define the size of an image by using real pixel values. In relative mode percents are used.

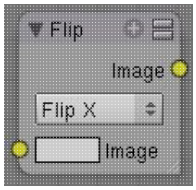
For instance X: 0.50 and Y: 0.50 would produce image which width and height would be half of what they used to be. When expanding an image greatly, you might want to blur it somewhat to remove the square corners that might result. Unless of course you want that effect; in which case, ignore what I just said.

Use this node to match image sizes. Most nodes produce an image that is the same size as the image input into their top image socket. So, if you want to uniformly combine two images of different size, you must scale the second to match the resolution of the first.

Flip Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Distort â Flip



Flip node
This node flips an image at defined axis that can be either X or Y. Also flipping can be done on both X and Y axis' simultaneously.

You can use this node to just flip or use it as a part of mirror setting. Mix half of the image to be mirrored with its flipped version to produce mirrored image.

Displace Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Distort â Displace

Ever look down the road on a hot summer day? See how the image is distorted by the hot air? That's because the light is being bent by the air; the air itself is acting like a lens. This fancy little node does the same thing; it moves an input image's pixels based on an input vector mask (the vector mask mimics the effect of the hot air).

This can be useful for a lot of things, like hot air distortion, quick-and-dirty refraction, compositing live footage behind refracting objects like looking through bent glass or glass blocks, and more! Remember what HAL saw in 2001:Space Odyssey; that distorted wide-angle lens? Yup, this node can take a flat image and apply a mask to produce that image.

The amount of displacement in the X and Y directions is determined by

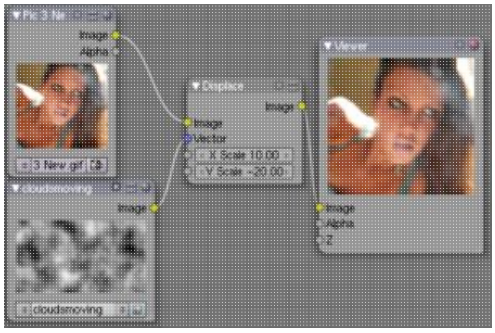
- The value of the mask's channels:
- The scaling of the mask's channels

The (red) channel 1's value determines displacement along the positive or negative X axis. The (green) channel 2's value determines displacement along the positive or negative Y axis.

If both the channels' values are equal (i.e. a greyscale image), the input image will be displaced equally in both X and Y directions, and also according to the X scale and Y scale buttons. These scale button act as multipliers to increase or decrease the strength of the displacement along their respective axes. They need to be set to non-zero values for the node to have any effect.

Because of this, you can use the displace node in two ways, with a greyscale mask (easy to paint, or take from a procedural texture), or with a vector channel or RGB image, such as a normal pass, which will displace the pixels based on the normal direction.

Example



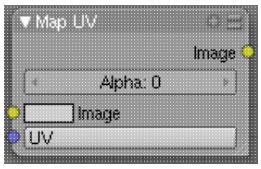
Music Video Distortion Example Using Displace
In this example, she's singing about dreams of the future. So, to represent this, we use a moving clouds texture (shot just by rendering the cloud texture on a moving plane) as the displacement map. Now, the colors in a black and white image go from zero (black) to one (white), which, if fed directly without scaling would only shift the pixels one position. So, we scale their effect in the X and Y direction. Upon reviewing it, sometimes stretching in both the X and Y direction made her face look fat, and we all can guess her reaction to looking fat on camera. SO, we scale it only half as much in the X so her face looks longer and thinner. Now, a single image does not do justice to the animation effect as the cloud moves, and this simple noodle does not reflect using blur and overlays to enahnce (and complicate) the effect, but this is the core.

Photos courtesy of Becca, no rights reserved. See also some movies of this node in action, made by the wizard programmer himself, by following this external link

Map UV Node

Panel: Node Editor → Node Composition

Menu: â § ShiftA â Distort â Map UV



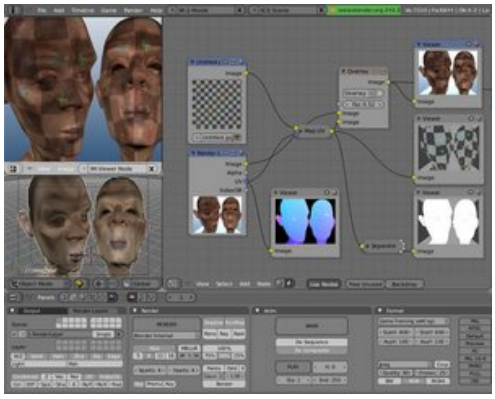
So, I think we all agree that the problem is...we just don't know what we want. The same is true for directors. Despite our best job texturing our models, in post production, inevitably the director changes their mind. "Man, I really wish he looked more ragged. Who did makeup, anyway?" comes the remark. While you can do quite a bit of coloring in post production, there are limits. Well, now this little node comes along and you have the power to **retexture your objects *after* they have been rendered**. Yes, you read that right; it's not a typo and I'm not crazy. At least, not today.

Using this node (and having saved the UV map in a multilayer OpenEXR format image sequence), you can apply new flat image textures to all objects (or individual objects if you used the very cool ID Mask Node to enumerate your objects) in the scene.

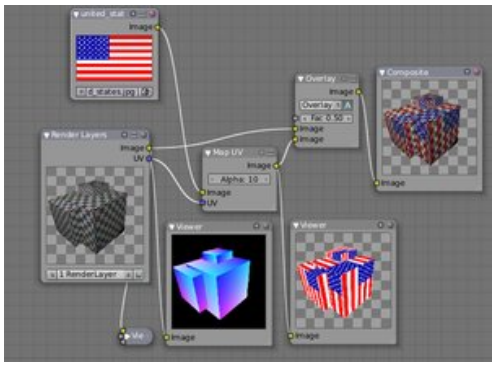
Thread the new UV Texture to the Image socket, and the UV Map from the rendered scene to the UV input socket. The resulting image is the input image texture distorted to match the UV coordinates. That image can then be overlay mixed with the original image to paint the texture on top of the original. Adjust alpha and the mix factor to control how much the new texture overlays the old.

Of course, when painting the new texture, it helps to have the UV maps for the original objects in the scene, so keep those UV texture outlines around even after all shooting is done.

Examples



Adding a Grid UV Textures for Motion Tracking
In the example to the right, we have overlaid a grid pattern on top of the two Emo heads after they have been rendered. During rendering, we enabled the UV layer in the RenderLayer tab (Buttons window, Render Context, RenderLayer tab). Using a mix node, we mix that new UV Texture over the original face. We can use this grid texture to help in any motion tracking that we need to do.



Adding UV Textures in Post-Production
In this example, we overlay a flag on top of a cubie-type thing, and we ensure that we Enable the Alpha premultiply button on the Mix node. The flag is used as additional UV Texture on top of the grid. Other examples include the possibility that we used an unauthorized product box during our initial animation, and we need to substitute in a different product sponsor after rendering.

Of course, this node does NOT give directors the power to rush pre-production rendering under the guise of "we'll fix it later", so maybe you don't want to tell them about this node. Let's keep it to ourselves for now.

Introduction

In addition to modelling and animation, Blender has a fully functional Video Sequence Editor as well as an advanced node-based editor that also manipulates a video stream. Compositing Nodes operate equally well on images or video streams, and can apply detailed image manipulation on the stream.

Operating at a higher conceptual level, and used later in the video production process, Blender's legacy Video Sequence Editor (VSE) operates on a set of entire strips at a time, as a chunk of footage. The many parts of Blender work together in typical work flow fashion:

1. Model to construct the objects
2. Assign Materials and introduce Lighting to color the objects
3. Animate your objects to make them move
4. Render layers of video using cameras
5. Use Compositing nodes to:
 1. Enhance the images by adjusting colors, adding in-scene special effects
 2. Layer the images into a composite image sequence (strip)
6. Assemble the video strips together to make a movie using the VSE.

The VSE within Blender is a complete video editing system that allows you to combine multiple video channels and add effects to them. Its functionality has been inside Blender since the beginning. Even though it has a limited number of operations, you can use these to create powerful video edits (especially when you combine it with the animation power of Blender!) Furthermore, it is extensible via a plugin system to perform an unlimited number of image manipulations.

Using the VSE, you load multiple video clips and lay them end-to-end (or in some cases, overlay them), inserting fades and transitions to link the end of one clip to the beginning of another. Finally, add an audio track so you can synchronize the timing of the video sequence to match it. The result of using the VSE is your finished movie.

FFMPEG Support
Support for exporting an avi/quicktime movie using FFMPEG does work, currently (since 2.44) only within the Linux and Windows builds. With FFMPEG support, you are able to save the audio track with your video.

In addition to being a work area for arranging video strips, the VSE workspace can show you different aspects of the composite result, for the current frame:

- Chroma: Color hue and saturation
- Luma: Brightness/contrast
- Image: Colors (what you see)
- Sequence: Video strips

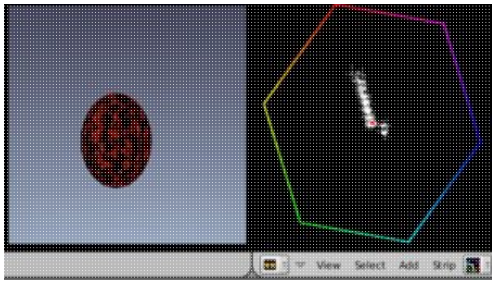


VSE Workspace/Image Preview Modes

In the Chroma, Luma, and Image modes, a channel selector appears; channel 0 is the result of compositing the strips with their special effects strips. Channel 1 is what the current frame's image from the strip in channel 1 looks like (channel 1 is at the bottom of the heap). The display of these modes is either the composite (channel 0) or the frame from the strip (channels 1 through n).

Zoom the view of any of these workspaces by scrolling your middle mouse wheel.

Chroma Vectorscope



Example VSE Chroma Preview
For the selected channel, this display shows the colorspace of the image inside a hexagon. Each point of the hexagon is a primary color: red, magenta, blue, cyan, green, and yellow. Black is at the center, and overall saturation is scaled as dots closer to the outside. The example to the right shows that the image has a lot of red (50% saturation) and small amount of blue, with no green.

Use this display to check for too much color saturation. While over-saturated images look great for op-art and computer displays, they stink when shown on the big screen TV. Use the Alt-Animation key to scrub the video; this display will update with a new/revised map for each frame. Just like watching the Image preview to see what it looks like, watch the Chroma Vectorscope to watch for color use.

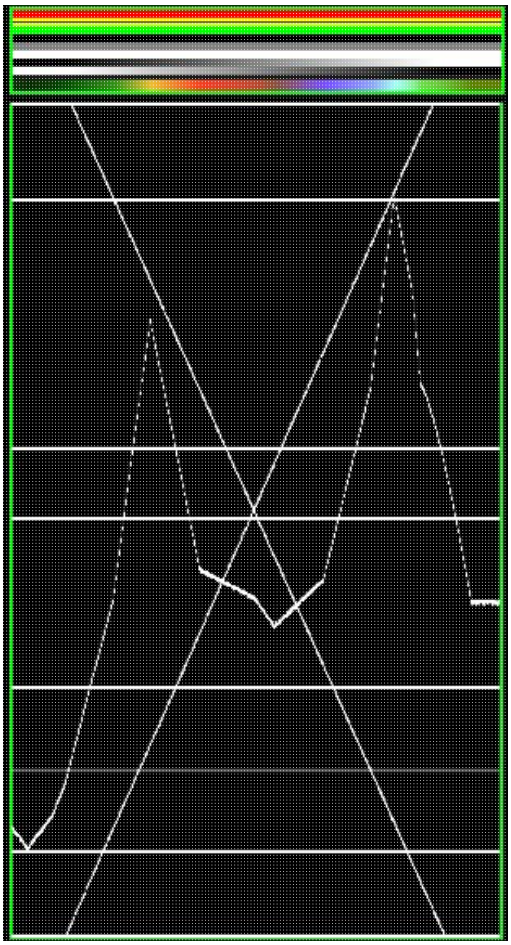
Luma Waveform

For the selected channel, brightness, or luminosity, is mapped with this display.

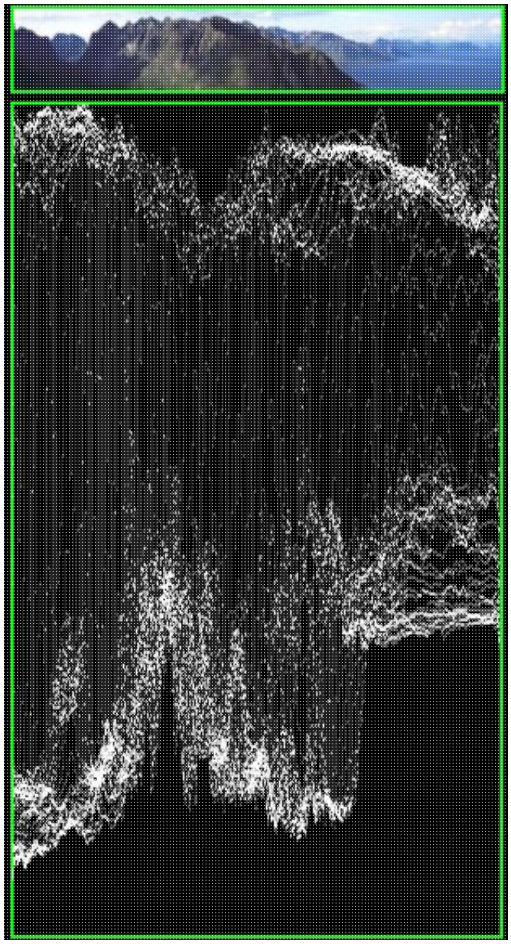
Note:
The original explanation seems â← to me â← quite confused, so I proposes here my own interpretation of this mode (deduced from my experiments)â←¦ If someone knows more about it, please correct it!

In this mode, the vertical axis represents the luminosity: null (full black) at the bottom, full (full white) at the top; the horizontal axis is a mapping from the horizontal axis of the frame. There are as much curves as lines in the frame: each one of this curves represents the luminosity of the pixels of one line. Moreover, the color of a pixel in this mode represents the number of pixels from the matching column of the frame sharing the same luminosity â← i.e. the number of curves that cross at this point (black/transparent, for no pixel, white/opaque for at least 3 (?) pixels).

It's not easy to explain, try it with various pictures (rather with natural pictures, the curves are more easy to see) and you might understand better. Look at the two examples below, too.



'Simple' picture.
The various horizontal lines in the Luma waveform match the uniform-color lines of the picture.
Note that the 'grey 20%' one-pixel width line (inside the yellow strip) is represented in the Luma waveform by a grey line.
The two lines drawing an 'X' are from the two linear tone shades (whiteâ black and blackâ white).
Finally, the broken line matches the complex tone shade at the bottom of the picture.



A 'real' picture.
The curves are quite visible.
We found a luma of 80-100% for the sky,
a luma around 40% for the sea,
and a luma of 10-20% for the mountains,
growing around 40% for the sunny part.

Examples of VSE Luma Previews.
Note that the pictures (first green frame, at the top) are only 50px high, to limit the number of curves displayed in the Luma waveform!

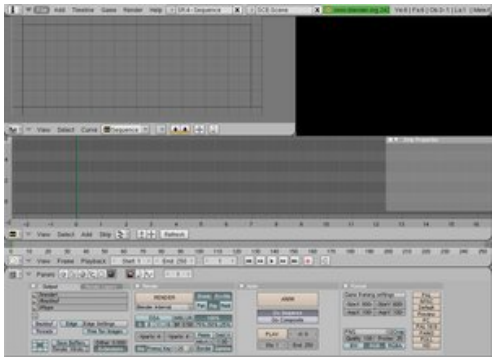
Use this display to check for appropriate contrast and luminosity across all frames in the channel. When spots in the film that should have even illumination don't, it looks like a flashbulb went off or an extra light was suddenly turned on. This can happen if two strips were rendered or shot under different lighting conditions but are supposed to be contiguous.

Image Preview

In the upper right window pane of the Sequence screen layout is another VSE window, this one set to Image Preview mode. It shows you what the resulting video will look like when saved.

Sequence Mode

The main working mode for adding strips and moving them around, cutting, grouping (making meta) and splicing them through special effects.



Default Sequence screen layout
Blender comes with a few screen layouts, one of which is 5-Sequence. This screen layout is shown to the right. The three main windows are an Ipo window and VSE window on the top row, and another VSE window in the middle, and a Timeline window below that, with the base Buttons windows at the bottom.

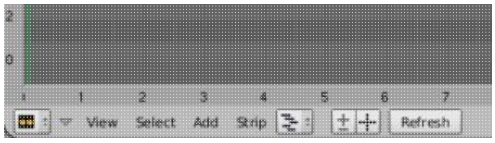
The Ipo Window is in Sequence mode; when you animate a sequence, it's Ipo curves will show up here, so you can make fine adjustments to its motion. The black VSE window is set to show you a preview of the finished sequence. Upon opening, it is black because there isn't anything to show.

The middle VSE window, in Sequence mode, is your work area. The VSE window is in the middle, because it is where you will be doing most of your work; being in the middle means your mouse has to travel the least amount of distance, and you are efficient. Keep this in mind when you are arranging *any* screen layout.

The Timeline window below it lets you select where (time-wise) in your animation you want to work, and allow you to quickly set the range of animation you want to scrub. Use the Timeline window's VCR buttons to playback the animation.

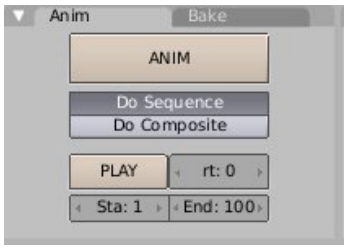
Scrub
Just like scrubbing a dirty pot with a brush, where you go back and forth, back and forth, until it is clean, the term *scrubbing* in video is going back and forth over a small range of frames, examining and correcting errors, until it is done.
Blender's Video Sequence Editor is a flexible workbench for editing your video footage. It is used to review your footage, and glue many sequences of your movie together. It offers a number of built-in and plug-in effects to transition from sequence to sequence, providing advanced hollywood-style effects for a professional looking video.

Overview of the Video Sequence Editor Window



Video Sequence Editor in Sequence display mode

The Video Sequence Editor has a header (where the menu and mode are shown) and a workspace, and works in one of four modes. In Sequence mode, the workspace is horizontally-striped into channels, and each video strip will go in a horizontal channel. Each channel is numbered on the left-hand side, starting from 0 (you can't put anything thing in this special one!) and going up; the picture to the right shows three channels (zero to two). Stripes toward the bottom are more dominant, which we'll get to in a minute. In the x direction, seconds of animation or frames of animation (T to choose) are used as the measure of time (seconds 1 through 7 are shown). You can scale the time using the zoom keys or mouse actions (see the Reference for more info).



Enable Sequence Output
When you click Render or Anim to generate an image or video, Blender has a choice of what image to compose for the current frame/scrub range:

- Current Scene layer result
- Sequence Editor channel 0 result
- Composition Node Editor renderlayer result

Tell Blender to use the output of the Video Sequence Editor by enabling "Do Sequence" in the Buttons Window, Scene Render buttons, Anim panel.

View Menu

As usual, this menu controls what and how you view in the workspace.

Mode: Sequence

Hotkey: â § ShiftSpace, CtrlUparrow, CtrlDownarrow

Menu: View -> Maximize Window

Use this when working arranging a lot of strips and you want to use all of your screen to work.

Mode: Sequence

Hotkey: T

Menu: View -> Show Frames, View -> Show Seconds

Time toggles the units of measure across the bottom of the workspace between seconds and frames. Seconds depends on the Frames/second setting in the Scene Render Format panel in the Buttons window.

Mode: Sequence

Menu: View -> Lock Time to Other Windows

Locking Time to other windows means that if you change your position in time (by left-clicking in the workspace to move your vertical green cursor), other windows will update as well to reflect what the video will look like at that time.

Mode: Sequence

Hotkey: . NumPad

Menu: View -> View Selected

Zooms in the display to fit only the selected strips

Mode: Sequence

Hotkey: â Home

Menu: View -> View All

Zooms (out) the display to show all strips

Mode: Sequence

Hotkey: â § ShiftAltA

Menu: View -> Play Back Animation in 3D View

Plays the animation for the selected scrub frame range in all windows: In VSE windows, shows you their respective display (Image, Chroma, Luma). In any 3D windows, shows you the objects moving.

Mode: Sequence, Image, Chroma, Luma

Hotkey: AltA

Menu: View -> Play Back Animation

Plays the animation for the selected scrub frame range in the window.

Select Menu

This menu helps you select strips.

Mode: Sequence

Hotkey: A

Menu: Select -> Select/Deselect All

Selecs all the strips loaded.

Mode: Sequence

Hotkey: B

Menu: Select -> Border Select

Begins the Box mode select process. Click and drag a rectangular lasso around a region of strips in your Sequence workspace. When you release the mouse button, the additional strips will be selected.

Add Menu

Mode: Sequence

Hotkey: Space

Menu: Add

This is the main menu you will be using. In general, you load up your strips, create strips of special transition effects, and then animate out your sequence by selecting "Do Sequence" and clicking the Anim button. You can use the Add menu in the header, or hover your mouse cursor over the Sequence workspace and press Space.

Clips can be Huge
A three minute quicktime .mov file can be 140Megs. Loading it, even over a high-speed LAN can take some time. Don't assume your computer or Blender has locked up if nothing happens for awhile.

Add Movie and Image (Stills or Sequences)

First, let's add a clip:

- A movie clip in the Audio-Video Interleaved format (*.avi file)
- A movie clip in the Apple QuickTime format (*.mov)
- A single still image to be repeated for a number of frames (*.jpg, *.png, etc.)
- A numbered sequence of images (*-0001.jpg, *-0002.jpg, *-0003.jpg, etc, of any image format)
- One or more images from a directory
- A Scene in your .blend file.

Blender does not care which of these you use; you can freely mix and match any of them. They all become a color-coded strip in the VSE:

- Blue is used for Avi/mov codec strips
- Grey is a single image that is repeated/copied
- Purple is an image sequences or group of images played one after the other
- Green is an Audio track

When you choose to add one of these, the VSE window will switch to a file browser for you to select what you want to add. Supported files have a little rectangle next to their name (blue for images, green for clips) as a visual cue that you can pick them successfully:

- When adding a Movie or Movie+Audio:

 LMB LEFT CLICK to put the name of the file into the text box at the top; this selects a **single** file (like a movie)

- In the case of (numbered) image **sequences**, you have a choice:

Directory: RMB right-click on a directory name, and all files in that directory will be brought in as part of the image, in sort order, one image per frame
Range: Navigate into the directory and right-click and drag over a range of names to highlight multiple files. You can page down and continue right-click-dragging to add more to the selection
Batch: Shift-right-click selected non-related stills for batch processing; each image will be one frame, in sort order, and can be a mix of file types (jpg, png, exr, etc.)
All: Press A to select/deselect All files in the directory.

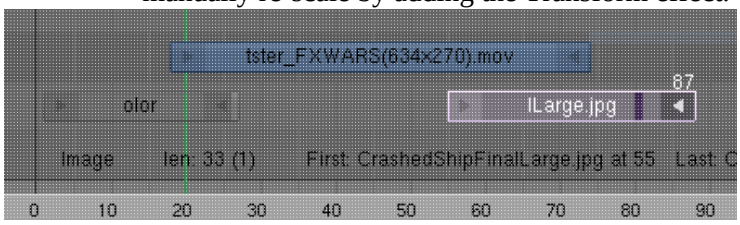
When you click the Select <whatever> button, the window pane will switch back to VSE, and the strip will be rubber-banded to your mouse. You cannot load multiple movies at the same time by right-clicking them; no movies load if you right click them. Right-clicking only works for images.

Error: The selected file is not a movie or FFMPEG support not compiled in!
means that the file is not a movie that Blender can recognize, or **you selected with the wrong button**. You get this error message because you *right*-clicked on a movie file, OR you don't have a codec that can decode the avi file. If it's the latter, find a codec so you can play the file outside of Blender, and then you will be able to load it. If it's the former, you must left-click to select movies.
In order to add items to the VSE, left-click for movies, left-click for single images, or right-click and drag for image sequences. Move your mouse to the frame/time and stripe you want, and click to break the rubberband and drop the strip in place (in a channel and starting at a frame).

When you add an image, Blender makes it into a 50-frame strip, which means that image will be in your video for two seconds (at 25 fps â← PAL). Aside from re-positioning it, you will want to scale it by RMB -clicking on either the start or end arrow, and dragging left or right. As you move, the frame number updates to say where the arrow is. Click LMB to validate, or RMB to cancel the modification.

Dealing with Different Sizes:

Dealing with different sized images and different sized outputs is tricky. Think like a pixel. If you have a mis-match between the size of the input image and the render output size, the VSE does try to auto-scale the image to fit it entirely in the output. This may result in clipping. If you do not want that, use Crop and/or Translate in the Input panel to move and select a region of the image within the output. When you use Crop or Translate, the auto-scaling will be disabled and you can manually re-scale by adding the Transform effect.



If you scroll up the workspace, you will see an information channel (at vertical location channel 0) that gives you some helpful hints about the active strip. The example below shows a color strip from frames 1 to 25, then a mov file, and then an image strip. The info channel shows handy information about the image strip, whose name has been scrunched in the strip display, but is clearly spelled out in the information strip.

9999 frames go by (IMAGE strips only!)
Ok, so that was a very obscure reference to a song about 99 balloons, but we really have not anticipated how fast Blender has moved into mainstream video editing. Unfortunately, we initially reserved 4 digits for the filename of each video image sequence set. While that provides for up to 400 seconds of video (about 5 minutes US), with Blender moving into movies, you need to break up IMAGE strips into 4 digits only, and others 5 digits (10000-19999), (20000-29999), etc. Important: that only affects IMAGE strips at the moment. All the other strip types work fine with up to 300,000 frames (approximately 3 hours, read: Ben Hur :)).
Codecs

You must have a codec on your machine that can decode the avi file. Blender does not control these. For example, the XviD codec is available from www.xvid.org

FFMPEG Support

If you are using a Blender build with FFMPEG support, you will be able to load audio and video strips together; select Movie+Audio(HD) and when you drop the strip, the strip will split into an audio and video channel strips.

Add Scene

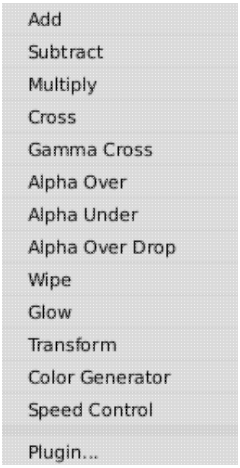
You can add the virtual image output of a Scene in your current .blend file as well. Select the scene from the popup list, and a strip will be added and rubberbanded to your mouse just like a movie or image. The strip length will be determined based on the animation settings in that scene (not the current scene, unless the VSE is operating in the same scene).

When adding a Scene strip, please note that, in order to show you the strip in the VSE Image preview mode, Blender must render the scene. This may take awhile if the scene is complex, so there may be a delay between the time you select the scene and the time the strip appears. To reduce the delay, simplify the scene rendering by selecting fewer layers to render.

Add Audio

The VSE can incorporate an audio channel which you can hear as you scrub. Add an audio track when you are trying to time your video/animation to an audio track, or vice versa. Please refer to the Audio Sequences section for more information.

Add Effect



Available Built-in Effects

Blender offers two categories of effects: Built-in and Plug-in. The built-in effects are listed to the right. They are built-in to Blender and everyone has them. The plug-in effects are separate files in a sequence-plugin directory on your PC that are loaded as they are needed. While a standard set of plugins are distributed when you installed Blender, everyone's computer may have a different set.

Mode: Sequence

Hotkey: Space

Menu: Add -> Effect

Every Built-in effect is explained later individually, but they all are added and controlled in the same way. To add an effect strip, select one base strip (image, movie, or scene) by RMB clicking on it. For some effects, like the Cross transition effect, you will need to â § ShiftRMB a second overlapping strip (it depends on the effect you want). Then select Add -> Effect and pick the effect you want from the pop-up menu. When you do, the Effect strip will be shown above the source strips. If it is an independent effect, like the color generator (described later), it will be rubberbanded to your mouse; click to drop the strip.

Since most Effects strips depend on one or two source strips, their frame location and duration depends on their source strips. Thus, you may not be able to move it; you have to move the source strips in order to affect the effect strip.

To use an effect that combines or makes a transition between (or composites) two strips, you must Box select or shift-right-click two of them. When you add the effect strip, it will be placed in a channel above the two in Grab mode (click to drop it on a channel). Its duration will be the overlap between the two strips as a maximum.

With some effects, like the AlphaOver, the order in which you select the strips is important. You can also use one effect strip as the input or source strip with another strip, thus layering effects on top of one another.

Note: The only exception is the Color Generator effect. It does not depend on a base strip; you can add and position it independent of any other strip. Change the length as you would any strip.

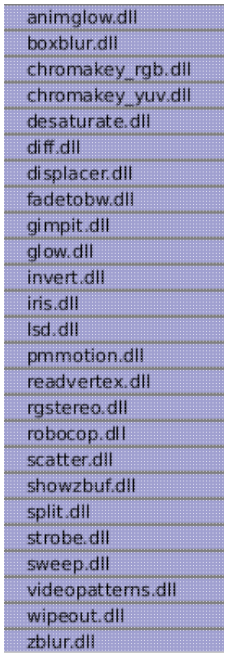
Mode: Sequence, Effects Strip Selected

Hotkey: C

Menu: Strip -> Change Effect

If you picked the wrong effect from the menu, you can always change it by selecting the strip (RMB) and using the Strip->Change Effect selection. Or, you can press Change to switch effects on a selected Effects strip.

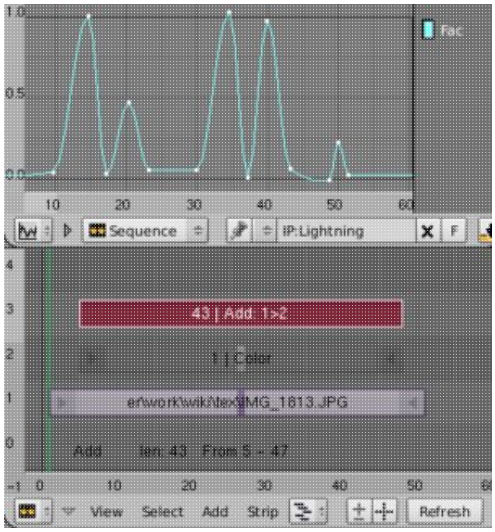
Adding Plugin Effects



Sequence Plugins are special little routines written by special programmers in the C language as a dynamic load library (.DLL). A DLL can be loaded at any time (dynamically) as it is needed, so it "plugs in" to Blender. (In case you wondered: the extension is platform dependent. These files are named .so (shared object) on Linux e.g.)

The image to the right shows the Sequence DLL's that I have available on my system. Each of them do some special effect indicated by their name or as explained on the Blender Resource Page for Plugins or programmer website. For example, the Iris plugin transitions between two strips by opening an expanding hole in the middle of the first and letting the second one show through, like an iris of a camera opening up. Some of these plugins can be five or more years old and still work very well; Blender tries to ensure backward compatibility, and they should work independent of output format or resolution (size).

Animating Audio/Video Effects



Animating the Add Effect

The degree to which some effects manipulate the image (called a Factor) can be controlled over time. For example, over the course of 100 frames, you can vary the factor from 0 to 1, then down to 0.5, then back up to 1.0. You do this by defining an Ipo Curve for that effect. Some effects (Add, Multiply) can be animated

through the Ipo curve; others (Color Generator, Glow) are controlled through the Properties panel.

To the right is an example of animating the Add effect to produce the effect that lightning has. To add control points in the Ipo Window, CtrlLMB click anywhere, and a Factor curve will be added. In this case, we have added gray to our image in a jagged Ipo set. When (or if you have already) gotten into nodes, the use of this Factor is the same in many nodes, especially the Mix node. By default, when you add a special effect, like a Cross, Blender provides a smooth transition to the new picture, from 0.0 to 1.0, over the duration of the frames (the sfx strip). In the case of an add or multiply or subtract effects, Blender add the effect immediately, and keeps it there; the factor jumps to 1.0 and stays at 1.0. Now, we know how much you all like to control things, so you have the Ipo Window that can operate on a strip, or Sequence. This is why there is an Ipo Window in the upper left-hand corner of the screen layout. In general you:

- Select the effect strip
- Decide whether you want the Ipo curve to be relative to the start of the first base strip, or locked to the global frame number. If you want it locked to the global frame counter, enable the IPO Frame locked button
- Ctrl-LMB click control points in the Ipo window (Sequence mode) to define the transition factor and timing. To make a smooth transition to the final effect, just define one point at the starting frame with a Fac of 0.0, and another point at the ending frame with a Fac of 1.0. You can define a jagged curve with lots of points to get a very dynamic effect.

Not all effects can be animated; The built-in Glow is fixed, but there is an AnimGlow plugin which can be animated. Add, Multiply, Divide, Subtract, and Transform effects are always animatable.

Animating an audio strip affects the volume of the strip in the resulting composite. Use IPO animation on an audio strip to fade in/out background music or to adjust volume levels. Layered/crossed audio strips are added together; the lower channel does not override and cut out higher channels. This makes Blender an audio mixer. By adding audio tracks and using the IPO curves to adjust each tracks' sound level, you have an automated dynamic multi-track audio mixer!

To shorten the transition time of any effects strip (such as Cross, Wipe, or Transform) that operates on two strips, just shorten the duration by which the two strips overlap. This will result in a change to the length of the effects strip.

If Ipo Frame locked is enabled, Blender will use the Fac value from frame x in the Ipo window to apply the Effect at frame x in the VSE window. If it is disabled, the Ipo curve is relative to the start of the first base strip. Generally, you want the transition effect to be relative to the strips they operate on, but in certain matching moves or transitions, you can lock it to a specific frame.

Strip Menu

This menu allows you to operate on video strip(s) as a whole.

Mode: Sequence

Hotkey: M, AltM

Menu: Strip -> Make Meta Strip, Separate Meta Strip

A Meta-Strip is a group of strips. Select all the strips you want to group, and Make them into one meta. The meta spans from the beginning of the first strip to the end of the last one, and condenses all channels into a single strip, just like doing a mixdown in audio software. Separating (ungrouping) them restores them to their relative positions and channels.

Mode: Sequence

Hotkey: X, â § ShiftD

Menu: Strip-> Delete, Duplicate

If you have added a strip by mistake or no longer want it, delete it by pressing X or using this menu option. Duplicate a strip to make an unlinked copy; drag it to a time and channel, and drop it by LMB click.

Mode: Sequence

Hotkey: K

Menu: Strip -> Cut at Current Frame

While splicing two strips happens just by placing them finish-to-start, cut a strip by pressing K to cut. At the selected frame for the selected strips, K cuts them in two. Use Kut to trim off roll-ups or lead-ins, or roll-downs or extra film shot ("C" was already taken for Change).

Note on the 'cut'
When you 'cut' a strip, you don't really make a cut like it was with the 'old editing' on real film. In fact, you make a copy of the strip: the end of the original one is 'winded' to the cut point, as with the beginning of the new copy.

For example, imagine that you have a strip of **50** frames, and that you want to delete the first ten ones. You have to go to the **11th** frame, and hit K; the cut 'divides' your strip in two parts. You now can select the first small part (frames **1** to **10**), and delete it hitting X.
You might think that you have really erased the frames **1** to **10**, but there are still there, 'winded', as in a film reel, under your frame **11**: you just have deleted one of the tow copies of your strip created by the 'cut'. And you can at any time get your 'lost' frames back (just RMB -click on the left arrow of the strip, then Grab it to the left to display the desired number of frames again (or to the right to 'hide' more frames â← this is another way to remove frames at the beginning/end of a strip!).

This is at the heart of nearly every editor solution, and that's quite handy!

Mode: Sequence

Hotkey: â § ShiftS

Menu: Strip -> Snap to Current Frame

Position your cursor (vertical green line) to the time you want. Snap to current frame to start a strip exactly at the beginning of the frame. If your Time display is in seconds, you can get to fractional parts of a second by zooming the display; you can get all the way down to an individual frame.

Mode: Sequence

Hotkey: G

Menu: Strip -> Grab

Moves the selected strip(s) in time or in channels. Move your mouse horizontally (left/right) to change the strip's position in time. Move vertically (up/down) to change channels. Moving a strip above (to a higher

channel) and over another strip means that it will not display during the overlap; lower channel strips are displayed "in front" of higher channels â← **WARNING: this is exactly the opposite that what happens with most of the other video editors!** To get them both to be seen, you have to mix them using the Add, Multiply, Subtract or other compositing effects.

Mode: Sequence

Hotkey: Y

Menu: Strip -> Separate Images to Strips

Converts the strip into multiple strips, one strip for each frame. Very useful for slide shows and other cases where you want to bring in a set on non-continuous images.

Strip Properties

Blender 2.46 expanded the Strip properties from one panel to an entirely new Render (F10) subcontext.

Blender 2.46+
Mode: Sequence, Effects Strip selected

Menu: Buttons Window, Render Context (F10), Sequence sub-context



Strip Properties Panels
The properties for the strip are examined and set in various panels in the Buttons Window, Render Context, Sequence sub-context. Until we get around to documenting all of these new features, please see the [Release Notes.]

- Input - where to pull images from
- Filter - Image pre-processing
- Proxy - Use representatives of the real image, for low-powered PCs
- Edit - change properties of the strip

The panels for each of these sets of options and controls are shown to the right

Input Strip Properties Panel

Controls the source of the strip. Fields include filename, current frame, and auto-crop and auto-translate features, as well as an offset for starting and ending the strip.

This is here you can edit/update the path of the file used by a strip. Very useful when you moved it one way or the other â← Â this avoid you deleting and re-creating the strip!

You have two text fields for path, the first being the path of the parent directory (Dir), and the second the file name itself.

Filter Strip Properties Panel

Enables you to quickly set common image pre-processing options.

Premul: Premultiply Alpha
Flip: X - flips (reverses) the image left-to-right, Y reverses top-to-bottom, and Time reverses strip image sequence
Use Color Balance provides three filters to adjust coloration: Lift, Gaussian, and one other. Each pass can have a positive, or inverted effect by clicking the appropriate button. Set the amount of the effect by setting the color swatch; white (RGB 1,1,1) has no effect.

Proxy Strip Properties Panel

A proxy is a smaller image (faster to load) that stands in for the main image. When you Rebuild proxy Blender computes small images (like thumbnails) for the big images and may take some time. After computing them, though, editing functions like scrubbing and scrolling and compositing functions like cross using these proxies is much faster but gives a low-res result. Disable proxies before final rendering.

Edit Strip Properties Panel

Name
: You can name your strips

Blend Mode
: By default, a strip Replaces the output image of any lower-level strips. However, many other blending modes are available based on the strip type. For example, Alpha-Over automatically overlays the image on top of a lower level strip. These autoblending modes obviates the need for separate effect strips. Blend percent controls how much of an (even over time) effect the strip exerts

Mute
: Hides the strip so that it does not participate in the final image computation

IPO Frame Locked
: Uses the Sequence IPO curve to determine the Blend percent.

Start
: changes the starting frame number of the strip, which is the same as grabbing and moving the strip. Tip: when you add a strip, I like to just drop it and then use this field to place it at the frame I want, rather that trying to drag and drop in exactly the right place.

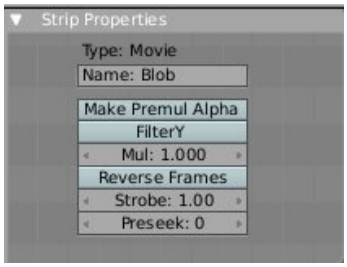
Chan

Changes the channel number of the strip, like G-Y.

Blender 2.45+
Mode: Sequence

Hotkey: N

Menu: Strip -> Strip Properties



Strip Properties
Press N to display a floating panel that shows you properties for the object selected; in this case a video Strip. By default, the name of a strip shown in the workspace is its filename. You can name strips by LMB clicking in the Name: field and entering a descriptive name; the workspace will display that name.

Use the Convert to Premul button if a strip has an Alpha (transparency) channel. Use FilterY if the strip is from broadcast video and has even or odd interlacing fields. Enhance the color saturation through the Multiply field. Play a strip backwards by enabling Reverse Frames. Tell Blender to display every nth frame by entering a Strobe value. Finally, when using MPEG video (VCD, DVD, XVid, DivX, â←¦), an image is built up over the course of a few frames; use the Preseek field to tell Blender to look backward and compose the image based on the n previous frames (e.g. **15** for Mpeg2 DVD).

For all effects, use the Strip Properties panel to control the effects strip; each effect has different controls, but they can all be set in the Properties panel. Control the length of the strip to vary the speed with which the transform happens. Regardless of whether they are built-in or plug-in, all effect strips do some special image manipulation, usually by operating on another strip or two in a different channel. The effect strip is shown in some channel, but its resultant effect shows up as Channel 0.

Working with Strips

Here are some common tasks that you will want to perform within the VSE.

Scrubbing

To move back and forth through your movie, use the Timeline window. LMB click and drag left/right in the timeline window, moving the vertical bar which indicates the current frame. As you do, the image for that frame is displayed in the VSE window.

Real-time scrubbing and image display is possible on reasonable computers when viewing an image sequence or movie (avi/mov) file. Scene images have to be rendered individually, which may take some time.

Selecting and Changing the Length of a Strip

To edit a movie or image strip:

- RMB in the middle of the strip selects the **entire** strip; holding it down (or pressing Grab) and then moving the mouse drags a strip around.
- RMB on the left arrow of the strip selects the **start** frame offset for that strip; holding it down (or pressing Grab and then moving the mouse left/right changes the start frame within the strip by the number of frames you move it:
 - ♦ If you have a 20-image sequence strip, and drag the left arrow to the right by 10 frames, the strip will start at image 11 (images 1 to 10 will be skipped). Use this to clip off a rollup or useless lead-in.
 - ♦ Dragging the left arrow left will create a lead-in (copies) of the first frame for as many frames as you drag it. Use this when you want some frames for transitions to the this clip.
- RMB on the right arrow of the strip selects the **end** frame of the strip; holding it down (or pressing Grab) and then moving the mouse changes the ending frame within the strip:
 - ♦ Dragging the right arrow to the left shortens the clip; any original images at the tail are ignored. Use this to quickly clip off a rolldown.
 - ♦ Dragging the right arrow right extends the clip. For movies and images sequences, more of the animation is used until exhausted. Extending a clip beyond its end results in Blender making a copy of the last image. Use this for transitions out of this clip.

Multiple selection
You can select several (handles of) strips by â § ShiftRMB -clicking: when you'll hit G, everything that's selected will move with your mouse â← this means that, for example, you can at the same time move a strip, shorten two others, and extend a forth one.

- STRIP EXTEND. With a number of Image strips selected, pressing E enters EXTEND mode. All selected strip handles to the "mouse side" of the current frame indicator will transform together, allowing you to essentially extend the strips that fall exactly on the current frame marker and having all others adjust to compensate.
- Additional Selection Methods
 - ♦ AltRMB and CtrlRMB on a strip will select the Left or Right handle of that strip and the neighboring handle on the next strip. Select with this method to move the boundary between two adjoining strips without affecting the outer limits.
 - ♦ AltCtrlRMB on a strip will select both handles of the strip, plus the neighboring handles on the immediately adjoining strips. Select with this method to move a strip that is between to others without affecting the selected strip's length.

Action Stops
When extending the start beyond the beginning or end after the ending, keep in mind that only the last image copies, so when viewed, action will stop on that frame. Start your transition (fade, cross) a little early while action is still happening so that the stop action is not that noticeable (unless, of course, you want it to be, like the 80's drama sitcoms).

Change the length of an effect strip by changing the start/end frame of the origin strips.

Zoom, Scale Display and Refresh Header Buttons



Zoom, Scale and Refresh buttons

These buttons are found throughout Blender on a window's header. LMB Click and drag the +/- button left and right to scale the display around the cursor (the vertical green line). Moving left zooms out, and right zooms in across frames. Moving up and down zooms in and out of channels.

Clicking the crosshair puts your mouse in a box-select mode. Select a region of the workspace by LMB clicking and dragging over a rectangular region of the workspace. When you release the mouse, the workspace display will zoom in to fit that region in your entire display.

Certain operations, like moving an object in 3D View, may not force the Sequencer to call for a refresh of the rendered image (since the movement may not affect the rendered image). If an image or video, used as a strip, is changed by some application outside of Blender, Blender has no real way of being notified from your operating system. To force Blender to re-read in files, and to force a re-render of the 3D View, click the Refresh button to force Blender to update and synchronize all cached images and compute the current frame.

Reverse Action

In the sequence buttons in the scene buttons (F10), select "flip time" in the "filters" panel.

Slow Motion, Fast Forward, Time Warp

Use the built-in Speed control strip documented on the next page.

Rendering a Video to an Image Sequence Set

In many cases, cuting and re-arranging (editing) a codec-encoded video strip will give you fits, because of the encoding algorithm that is used internally to reconstruct each image gets 'off' by a frame or two or three. To work directly on the 'raw' frame set, a very common technique is to import your video as a strip and render it out to series of individual frames, where each frame is stored in its own image file (JPG most commonly). To do so, Add->Movie and load your original video. Set your Format SizeX and SizeY (either to match the original, or different if you want to distort or upscale/downscale the video), set image type to JPEG, adjust your Quality settings, and in the Anim panel set your End: to the number of actual frames in the video strip. Click ANIM and a series of number files will be output to the top filespec in the Output panel. You can now delete the video strip, and Add->Image instead, and right click on the directory name to pull in all of the images, in sequence, that are within that directory. Now, when you cut at frame 4321, for example, the next frame of the second strip will *really* start with frame 4322.

Rendering a Video from an Image Sequence Set

Ridiculously easy (when you learned where the buttons are):

1. Add the sequence of images as described above
2. Set your Output file path and name to wherever you want to save the movie file (e.g. C:\My Documents\MyMovie) in the upper output box of the render buttons.
3. Change your Format to a movie file format (AVI, MOV, FFMPEG) and CODEC
4. Set your framerate to match whatever framerate the sequence is to be played back in. Under the Anim/Playback buttons.
5. Set your ANIM End: to the number of images in the sequence, and
6. ANIM

The single movie file is created and saved; the name is what you specified but with the starting frame and ending frame numbers appended (e.g. MyMovie0000-0250.avi)

Blender offers over a dozen effects that are built into Blender, and are therefore universal. Some operate on two strips; some on one, and some create a new strip. Each effect enhances your content in some way or allows professional-quality transitions.

Add Effect

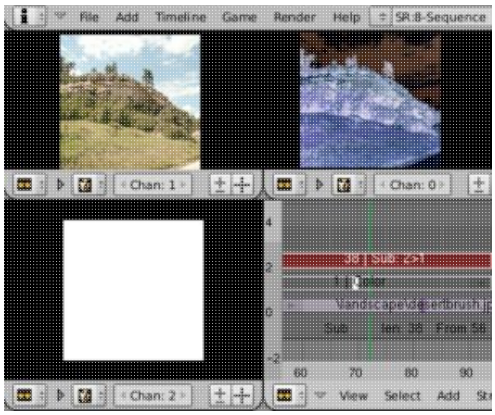


Can you hear the thunder?
The Add effect adds two colours together. Red and Cyan (Green and Blue) make White. Red and Blue make "Magenta" (i.e. Purple!). Red and Green make Yellow.

The Add Effect adds the colors of two strips together, Use this effect with a base image strip, and a modifier strip. The modifier strip is either a solid color or a black-and-whte mask, or another image entirely. The example to the right shows what happens when you add gray to an image, and animate the effect over time. The image gets bright because we are adding gray (R:.5, G:.5, B:.5) to say, a blue color (R.1, G:.1, B:.5) resulting in (R:.6, G:.6, B:1.0) which retains the original hue (relationship between the colors) but is much brighter (has a higher value). When applied to the whole image like this, the whole image seems to flash.

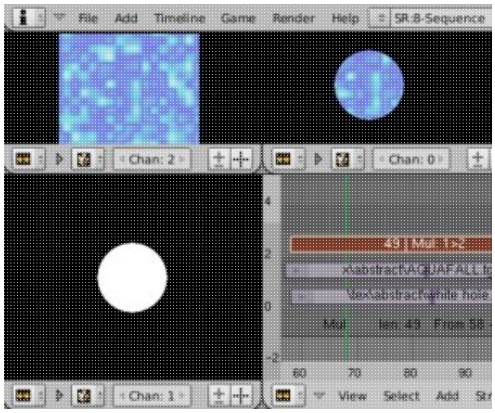
You can use this effect to increase the brightness of an image, or if you use a BW mask, selectively increase the brightness of certain areas of the image. The Mix node, in Add mode, does exactly the same thing as the Add sfx strip here, and is controlled the same way by feeding the Factor input.

Subtract Effect



Subtract Effect
This effect takes away one strip's color from the second. Make a negative of an image using this effect, or switch the order of the strips and just darken the strip. Subtracting a hue of blue from a white image will make it yellow, since red and green make yellow.

Multiply Built-in Effect



Multiply Effect.
The Multiply effect multiplies two colours. Blender uses values between **0.0** and **1.0** for the colours, he doesn't have to normalise this operation, the multiplication of two terms between **0.0** and **1.0** always gives a result between **0.0** and **1.0** (with the 'traditional' representation with three bytes â← like RGB(**124**, **255**, **56**) â← , the multiplications give far too high results â← like RGB(**7316**, **46410**, **1848**) â← , that have to be 'brought back', normalised â← just by dividing them by **256**! â← to 'go back' to range of **0** to **255**â←¦).

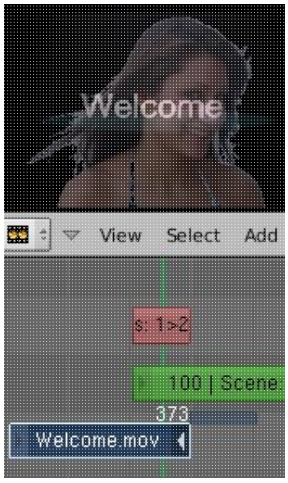
This effect has two main usages:

With a mask
: A mask is a B&W picture witch, after multiplication with a 'normal' image, only show this one in the white areas of the mask (everything else is black). The opening title sequence to James Bond movies, where the camera is looking down the barrel of a gun at James, is a good example of this effect.

With uniform colours
: Multiplying a colour with a 'normal' image allows you to soften some hues of this one (and so â← symmetrically â← to enhance the others). For example, if you have a brown pixel RGB(**0.50**, **0.29**, **0.05**), and you multiply it with a cyan filter (uniform colour RGB(**0.0**, **1.0**, **1.0**), you'll get a colour RGB(**0.0**, **0.29**, **0.5**). Visually, the result is to kill the reds and bring up (by 'symmetry' â← the real values remain unchanged!) the blues an greens. Physically, it is the same effect as shining a cyan light onto a chocolate bar. Emotionally, vegetation becomes more lush, water becomes more Caribbean and inviting, skies become friendlier.

Note
This effect reduces the global luminosity of the picture (the result will always be smaller than the smallest operand). If one of the image is all white, the result is the other picture; if one of the image is all black, the result is all black!

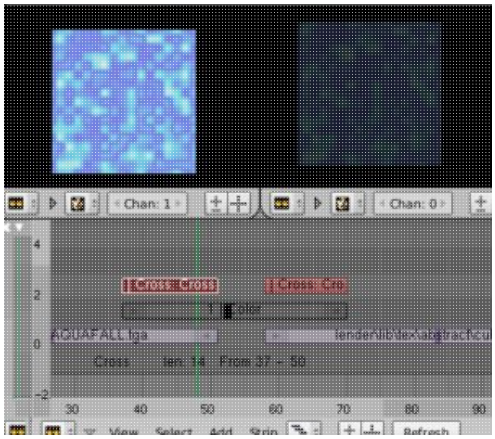
Cross and Gamma Cross



This effect fades from one strip to another, based on how many frames the two strips overlap. This is a very useful strip that blends the whole image from one to the other.

Gamma Cross uses color correction in doing the fade, resulting in a smooth transition that is easier on the eye.

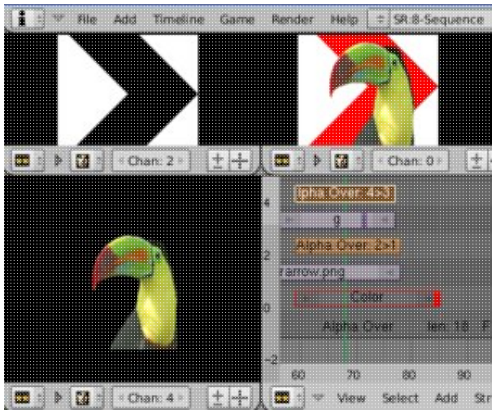
Fade to Black



Cross-Fade between Black
Many scenes fade to black, and then fade in from black, rather than directly from one to the other.

The strip setup to do this is shown to the right. The two strips are on Channel 1, and you Add->Color Generator strip to Channel 2, straddling the two main strips. Change the color to black, and add two Cross Effects; the first from Channel 1 to Channel 2 (black), and the second from Channel 2 to Channel 1. The first strip will fade to black, and then the second will fade in from black. Of course, you can use any transition color you want. Black is a relaxing intermediary; red is alarming. Use the dominant color in the second strip to introduce the second strip.

Alpha Over, Under, and Over Drop Built-in Effects



AlphaOver Effect

Using the alpha (transparency channel), this effect composites a result based on transparent areas of the dominant image. If you use a Scene strip, the areas of the image where there isn't anything solid are transparent; they have an alpha value of 0. If you use a movie strip, that movie has an alpha value of 1 (completely opaque).

So, you can use the Alpha Over/Alpha Under effect to composite the CGI Scene on top of your movie. The result is your model doing whatever as if it was part of the movie. The IPO Fac curve controls how much the foreground is mixed over the background, fading in the foreground on top of the background. The colors of transparent foreground image areas is ignored and does not change the color of the background.

Select two strips (â § ShiftRMB):

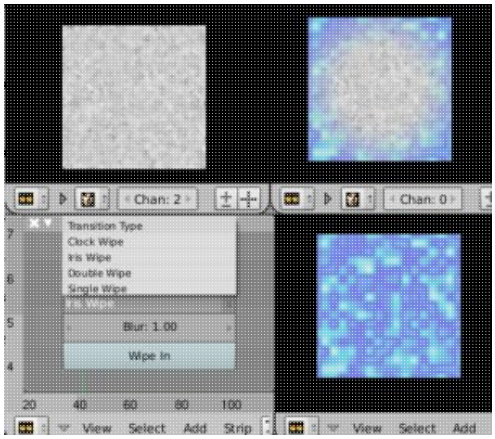
- With Alpha Over, the strips are layered up in the order selected; the first strip selected is the background, and the second one goes *over* the first one selected. The Factor controls *the transparency of the foreground*, i.e. a Fac of **0.0** will only show the background, and a Fac of **1.0** will completely override the background with the foreground (except in the transparent areas of this one, of course!)
- With Alpha Under, this is the contrary: the first strip selected is the foreground, and the second one, the background. Moreover, the Factor controls *the transparency of the background*, i.e. a Fac of **0.0** will only show the foreground (the background is completely transparent), and a Fac of **1.0** will give the same results as with Alpha Over.
- Alpha Over Drop is between the two others: as with Alpha Under, the first strip selected will be the foreground, but as with Alpha Over, the Factor controls the transparency of this foreground.

The example shows layering of AlphaOver effects. The very bottom channel is red, and an arrow is on top of that. Those two are AlphaOver to Channel 3. My favorite toucan is Channel 4, and Channel 5 alphaovers the toucan on top of the composited red arrow. The last effect added is tied to Channel 0 which will be rendered.

By clicking the PreMult Alpha button in the properties panel of the foreground strip, the Alpha values of the two strips are not multiplied or added together. Use this effect when adding a foreground strip that has a variable alpha channel (some opaque areas, some transparent, some in between) over a strip that has a flat opaque (Alpha=1.0 or greater) channel. If you notice a glow around your foreground objects, or strange transparent areas of your foreground object when using AlphaOver, enable PreMultiply. The AlphaOver Drop effect is much like the Cross, but puts preference to the top or second image, giving more of a gradual overlay effect than a blend like the Cross does. Of course, all of the Alpha effects respect the alpha (transparency) channel, whereas Cross does not.

The degree of Alpha applied, and thus color mixing, can be controlled by an IPO curve. Creating a Sine wave could have the effect of the foreground fading in and out.

Wipe Built-in Effect

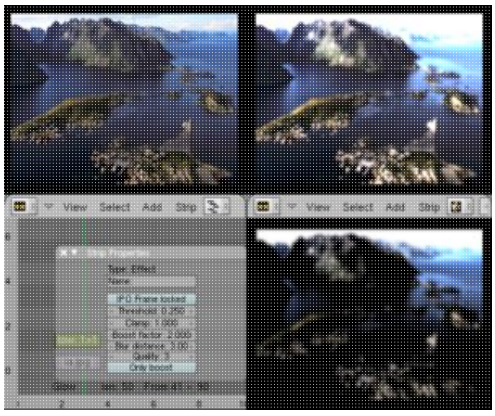


VSE Wipe Built-in Effect
Wipe transitions from one strip to another. This very flexible effect has four transition types:

- Clock: like the hands of an analog clock, it sweeps clockwise or (if Wipe In is enabled) counterclockwise from the 9:00 position. As it sweeps, it reveals the next strip.
- Iris: like the iris of a camera or eye, it reveals the next strip through an expanding (or contracting) circle. You can blur the transition, so it looks like ink bleeding through a paper.
- Double Wipe: Starts in the middle and wipes outward, revealing the next strip. It can also Wipe In, which means it starts at the outside and works its way toward the middle. You can angle and blur the wipe direction as well.
- Single Wipe: Reveals the next strip by uncovering it. Controls include an angle control so you can start at a corner or side, and blur the transition.

Note: some older plugins contain similar functionality.

Glow Built-in Effect



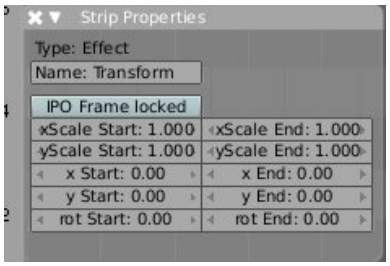
Example of a Glow effect applied to a picture.
Top left: base picture (Lofoten Islands, Norway â← source: wikipedia.fr);
Top right: result of the effect;
Bottom left: effect settings;
Bottom right: result with the Only boost button activated.
This effect makes parts of an image glow brighter by working on the luminance channel of an image. The Glow is the superposition of the base image and a modified version, where some areas (brighter than the

Threshold:) are blurred. With the Glow strip properties, you control this Threshold:, the maximum luminosity that can be added (Clamp:), a Boost factor: for it, the size of the blur (Blur distance:), and its Quality:. The Only boost button allows you to only show/use the 'modified' version of the image, without the base one. To "animate" the glow effect, mix it with the base image using the Gamma Cross effect, crossing from the base image to the glowing one.

Transform Built-in Effect



Transform is a swiss-army knife of image manipulation. It scales, shifts, and rotates the images within a strip. The example to the right shows what can be done with a single image. To make a smooth transition to the final effect, enable the IPO Frame locked button and define a curve in the Ipo Window (Sequence mode).



With the Transform strip selected, uses the properties panel to adjust the settings of this effect:

(x,y)Scale (Start,End):
: To adjust the scale (size). xScale Start defines the start width, xScale End the end width, yScale Start the start height, and yScale End the end height. The values higher than **1.0** will scale up the picture, while values lower than **1.0** will scale it down.

(x,y) (Start,End):
: To adjust the position (shifting). x Start defines the horizontal start position, x End, the end one; positive values shift the image to the right, negative values, to the left. y Start defines the vertical start position, y End, the end one; positive values shift the picture to the top, negative values, to the bottom.

rot (Start,End):
: The rotation is in degrees (**360** for a full turn) and is counter-clockwise. To make an image spin clockwise, make the end value lower than the start one (e.g. start it at 360 and go down from there).

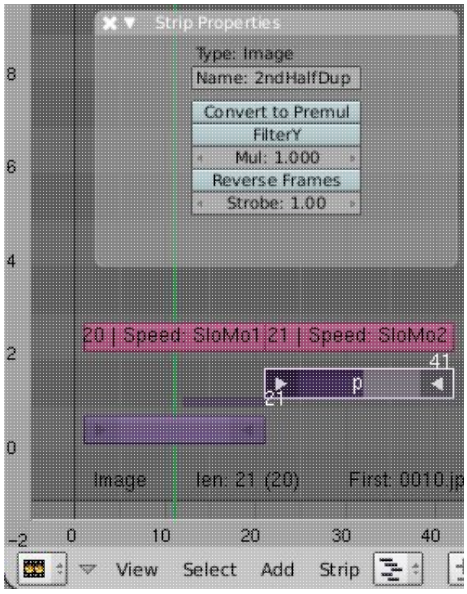
Color Generator Built-in Effect

This effect works by itself to create a color strip. By default, when it is created, it is 50 frames long, but you can extend it by grabbing and moving one of the ends. Click on the color swatch in the Effect panel under Sequencer buttons, which is under the Scene (F10) tab, to pick a different color (by default, it is gray). Use this strip crossed with your main movie to provide a fade-in or fade-out.

Speed Control Built-in Effect

Speed Control time-warps the strip, making it play faster or slower than it normally would. A Global Speed less than 1.0 makes the strip play slower; greater than 1.0 makes it play faster. Playing faster means that some frames are skipped, and the strip will run out of frames before the end frame. When the strip runs out of frames to display, it will just keep repeating the last one; action will appear to freeze. To avoid this, position the next strip under the original at a point where you want motion to continue.

Using Speed for a Slow-Motion Effect



50% Slow motion using Speed Control

Suppose you want to ssslooow your strip dowwwwwn. You need to affect the speed of the video clip without affecting the overall frame rate. Select the clip and Add->Effect->Speed Control effect strip. Click to drop it and press N to get the Properties. Set the Global Speed to be the factor by which you want to adjust the speed. To cut the displayed speed by 50%, enter 0.50. Now, a 30-frame clip will play at half speed, and thus display only the first 15 frames.

If you want the remaining frames to show in slo-mo after the first set is displayed, Kcut the strip in two, offset the second part (because the first slow-mo will actually run for more time/frames than shown), and add another Speed control as shown to the right. When the strip on Channel 1, as modified by the sfx strip on Channel 3 finishes, the strip selected on channel 2 starts, as modified by its sfx speed control in channel 3. The trick in creating the second strip is to

- Select the original strip and â § ShiftDuplicate it. In this case, it is a 20-frame image set.
- Drag the clone to half-way above the original and drop it.
- Select the left handle (start), grab it, and move it over half-way. This changes the start frame offset; the strip will now start playing at frame 11 (relative to within itself)
- Select the right handle (end), grab it and move it over half-again as much (in this case, 10 frames) This sets the duration to be 20 frames within your video.
- Add a 50% Speed Control for the first strip, and a 100% speed control for the second. For the second strip, you are already telling blender to play 10 frames over a 20-frame duration, which is already half speed, so you don't need the speed control to slow it down anymore, just to scale the frames selected to the duration of the clip.

That's it! Set your render to animate (in this example) all 40 frames.

Why not just extend the original clip out, you ask? Well, the Speed Control operates based on the number of frames it is going to show, divided by the number of frames it has to show them over. For example, if a strip has 20 frames in it, and you specify a speed factor of 50%, it knows it has to display 10 frames. If that strip is, say, stretched to cover 30 frames of video, it will play each of the 10 frames for 3 frames (3 x 10=30). Understanding the math will help you be effective and get the control to do what you want it to.

In very simple cases it will also work to just extend the original clip. This is a very nice feature intended for story boarding and for filling gaps. Add a speed control, ignore the IPOs and all the other parameters and just extend the right handle of the original clip. Remember that this is only possible, if the underlying clip ends at the right handle position. (If this isn't the case because you used the knife tool to create it, add a meta for the original clip first.).

Using speed control for frame matching

To get even finer control over your clip timing, you can use IPOs! This can be done in several ways.

- In the default configuration, your IPO curve simply map input to output frame scaled down to the range of [0-100] [0-1] in IPO curve coordinates. This is enough, if you want to add simple flickering effects.

The speed effect can do even more for you. You can do actual frame matching! That means: you have a certain position at which a certain frame should get displayed (maybe when matching your video track to an audio track) and blender will adjust the speed of the clip smoothly to make this happen.

- You can disable scaling in X direction by clicking on "IPO-frame locking". That means: input frame numbers are identical to IPO curve X-coordinates.
- You can disable scaling in Y direction by disabling "Scaling [0-1]". That means: output frame numbers are identical to the IPO curve Y-coordinates.

After that, use the N-keys window fo the IPO-window to place your control points in a frame exact way.

Changing video frame rates

You can use the speed control to change the frames per second (fps), or framerate, of a video. If you are rendering your video to a sequence set, you can effectively increase or decrease the number of individual image files created, by using a Global Speed value less than or greater than one, respectively. For example, if you captured a five-minute video at 30 fps and wanted to transfer that to film, which runs at 24 fps, you would enter a Global Speed of 30/24, or 1.25 (and Enable Frame Blending to give that film blur feel). Instead of producing 5*60*30=9000 frames, Blender would produce 9000/1.25=7200=5*60*24 frames. In this case, you set a Sta:1 and End:7200, set your Format output to Jpeg, 30fps, and image files 0001.jpg through 7200.jpg would be rendered out, but those images 'cover' the entire 9000 frames. The image file 7200.jpg is the same a frame 9000. When you read those images back into your film .blend at 24 fps, the strip will last exactly 5 minutes.

Plug-ins are just as easy to use as built-ins; they just require an extra step in finding them on your system and selecting the one you want based on its filename. By default, they are located in a Plugins directory within your Blender installation directory. In a default Windows installation, it is "C:/Program Files/Blender Foundation/Blender/plugins/", or wherever the path is entered within the User Preferences File Paths window.

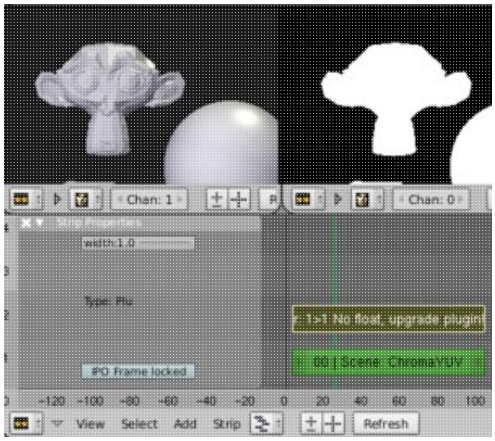
You find these DLLs by:

- Visit the Blender Resource Page for Plugins

- Hunt the internet and visiting programmer's websites
- Explore the plugins that are installed with Blender.

When you find them, download the version that is appropriate for your operating system (there are different flavors for Unix, Windows, and OS X), saving them in the Blender Plugin Sequence directory. The location of this directory on your machine is under your install directory or overridden by a setting in your User Preferences panel.

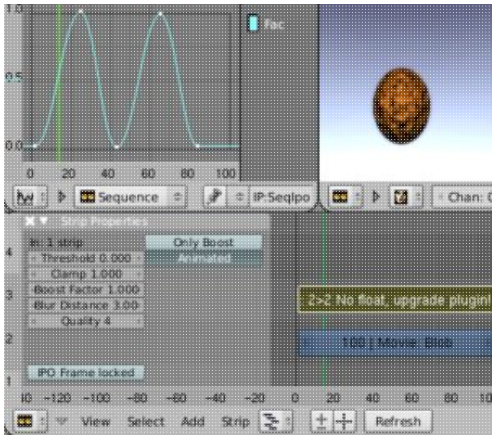
Alpha Matte Plugin Effect



VSE AlphaMatte Plugin
The Alpha Matte plugin converts the alpha channel of a strip to a black and white image corresponding to the alpha values; transparent is alpha 0 which is shown as black. Use this matte as a garbage matte, or as a multiplier against the selected strip so that a background strip will show through.

Note: Only your Scene and certain image types (PNG, EXR) support an alpha channel.

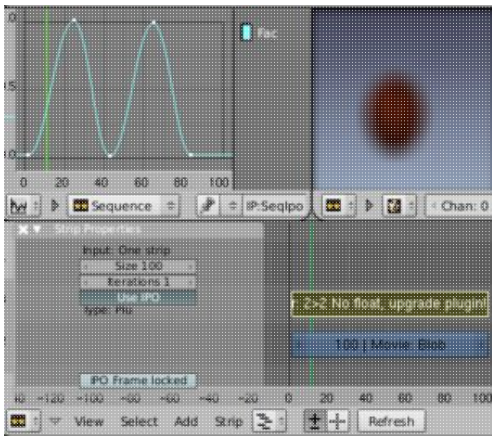
Animated Glow Plugin Effect



VSE AnimGlow Plugin
The Animated Glow plugin selectively increases the brightness of an image. Pixels with a brightness below the Threshold are not affected, and those selected will not be brightened above a Clamp value. The amount of glow is determined by the BoostFactor, and the Distance the glow extends outward is also selectable, as well as the Quality. In addition, you can see the composite glow, or just the amount of glow Boost that is added.

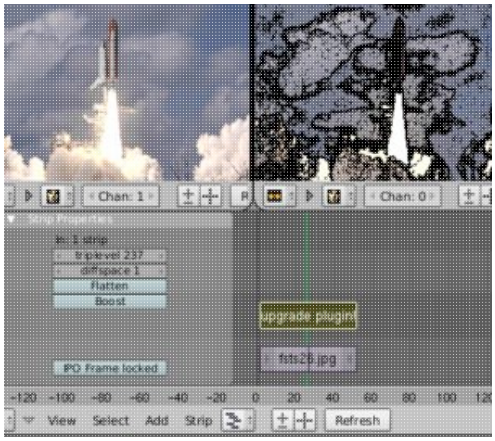
This example shows a glow in process at about 50% of a 100% Boost Factor; the image will glow twice over 100 frames, simulating a light being turned up and down.

Box Blur Plugin Effect (Pre 2.48 only; crashes 2.48+)



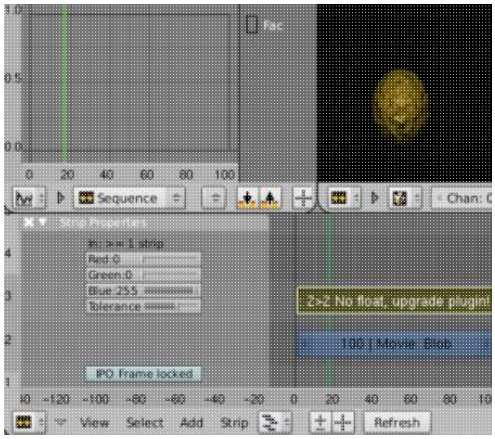
VSE BoxBlur Plugin
The BoxBlur plugin blurs the pixels in an image by the number of pixels defined by Size. The smoothness of the blur is increased by increasing Iterations. The blur can be animated over time by enabling Use IPO and creating an animation curve as shown. This example shows a blur in process at about 25% of a 100 pixels; the image will go in and out of blur twice over 100 frames, simulating a camera focusing.

Cartoon Plugin Effect



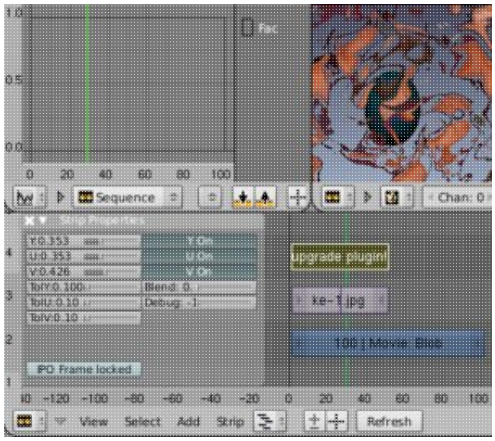
VSE Cartoon Plugin
The Cartoon plugin pixelizes and flattens the colors in an image, making a "real" image into something that would be drawn in a comic book. The trip level and difference space determines where the black "lines" should be drawn. Flatten and/or Boost image colors to add false color.

Chromakey RGB Plugin Effect



VSE Chromakey RGB Plugin.
This plugin turns the pixels which have the color determined by the RGB sliders, within an error Tolerance, transparent. The example above takes our blob on a shaded blue background and treats the background as a "blue screen", turning it all transparent. The same could have been done for green, which is called a "green screen". This allows us to now overlay the foreground blob over a new background using the builtin AlphaOver effect. This effect is not animatable.

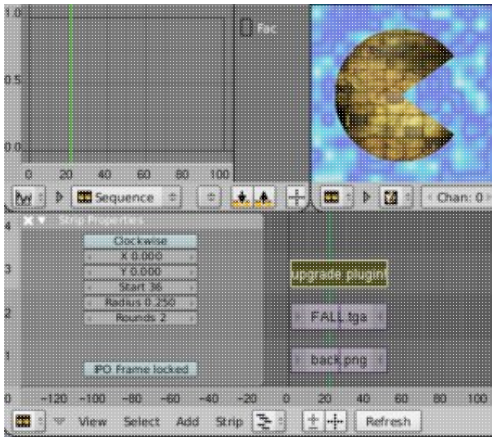
Chromakey YUV Plugin Effect



VSE Chromakey YUV Plugin.
Based on the luminance (Y) and chromiance (U and V) values of a foreground strip, this plugin shows the pixels of a background strip underneath the foreground. This plugin turns the pixels which have the color determined by the YUV value sliders, within an error tolerance defined by TolYUV sliders, transparent. The amount of blending is controlled by the Blend slider. The example above takes a swirly foreground image and puts our blob on a shaded blue background as the background. It treats certain colors within the foreground as the "blue screen".

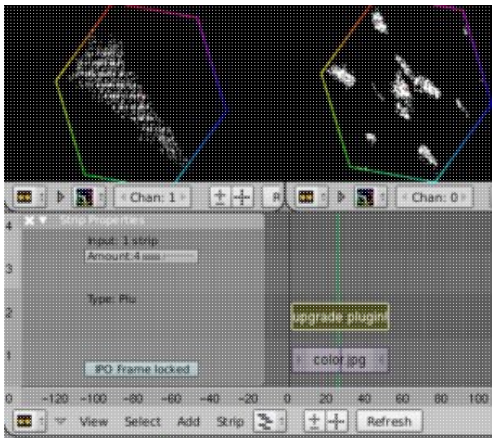
YUV chromakeying is often used when an actor is filmed in front of a black screen, and the YUV keying turns the black background to alpha, letting the background show through.

Clockwipe Plugin Effect



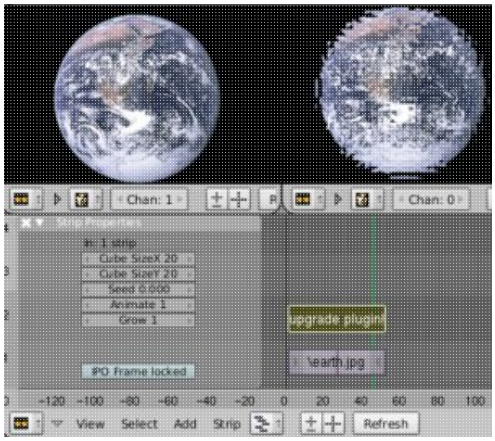
VSE Clockwipe Plugin
Remember the old countdown header in black and white movies? Based on the Starting angle, this plugin reveals the background image by wiping in a clockwise or counterclockwise direction. The Radius is how much of the image to reveal, and the clock goes round and round the number of times specified. As it completes one round, it reverts to the foreground image. Use this effect as an inventive wipe between scenes (using a radius of 0.5).

Desaturate Plugin Effect



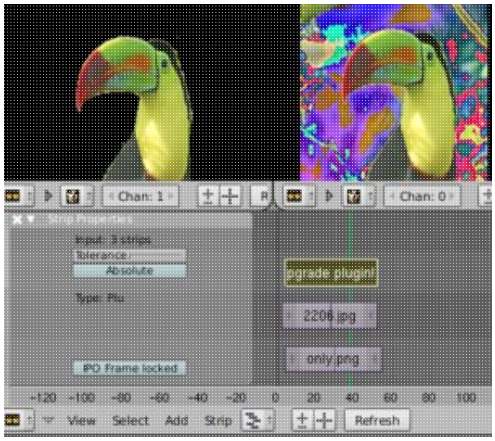
VSE Desaturate Plugin
This interesting effect removes the saturation from an image starting from the middle and working its way outward. The net effect is that for middle-range values, the colors in the image are made extreme. A saturation of zero is a black and white image, but in the middle, colors above the threshold become oversaturated (posterized). Notice that here we use the Chroma Vectorscope to see the color saturation map.

Dice Plugin Effect



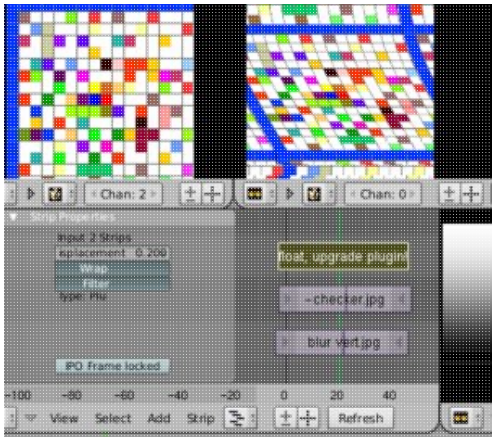
VSE Dice Plugin
Better than a Ginsu knife, this plugin slices and dices your image into cubes, and flips them. When animated and allowed to Grow (equal to 1), the effect grows over time to the cube size specified. A very interesting wipe effect, and you will never need a bandaid.

Diff Plugin Effect



VSE Difference Plugin
The Difference plugin subtracts one image from another, and shows you any difference between them, either as a color map (using color math) or as an Absolute black and white mask. If there is no difference, it passes through the original. Can you figure out what image we compared the Toucan to? In this example I played a little trick on you. I compared the image of the toucan only to the image of toucan in its environment. As a result, the environment is a color negative, while the toucan is the toucan since there is no difference in the toucan area between the two images. Gotcha!

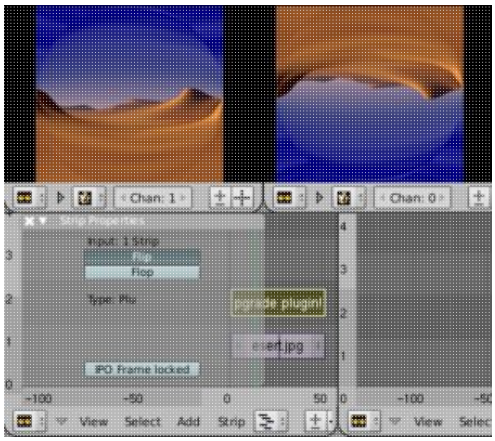
Displace Plugin Effect



VSE Displace Plugin
First selecting a vector map (a black and white image that tells the plugin what to do) and an image to manipulate, the Displace plugin shifts the pixels sideways and vertically. Use Wrap to make a wrap-around transition so that pixels shifted off the image on one side are used to fill in the other side. Filter does some anti-aliasing, making the image smoother.

A positive displacement shifts the pixels down and to the right, while a negative displacement shifts them up. Partially shown in the image to the right is the vector mask used for this example: a simple gradient that goes vertically from black to white. The shade of gray is multiplied by the displacement to determine how far to shift the pixels.

FlipFlop and FlipIt Plugin Effects



VSE FlipFlop Plugin
Flip an image upside down, or Flop it sideways left-right. Simple and cute.

FlipIt adds the option to rotate the image 90 degrees.

GreenScreen Plugin Effect



VSE GreenScreen Plugin

Combining live action footage, filmed in front of a green screen, and composited over a background, is simple and fast thanks to this amazing new (Feb 2007) plugin from Paprmh, proving once again that the Sequence plugins are alive and well.

Simple to use? Just load your greenscreen movie and load your background strip. Shift-select the greenscreen, then the background, and add this effect. When you click the check/ibuf2 button twice (until it turns yellow), the plugin automatically substitutes your background for the green removed. While nodes may be more flexible, they are about 15 times slower; using this plugin you can process a ten-minute run-time green-screen footage clip in a half an hour (on my Centrino(r) laptop).

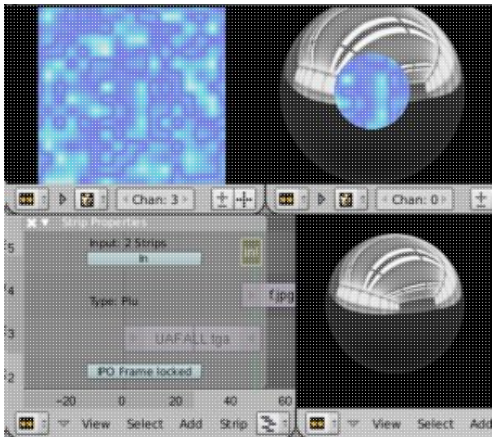
This easy-to-use plugin takes out the green with values between the Ymin and Ymax values, and leaves U (blue) and V (red) alone based on clipping values set with the sliders. It also has special filters for DV 4:1:1 footage that smooths out the color channels. The variable color spill factor, spill/+lum button get rid of that dark outline and any green that may have spilled onto the actors. The edge blur melds (anti-aliases) the edges of the keyed image to the background.

Download this plugin and visit his site for a complete explanation and tutorial that even covers garbage matte techniques using the Sequence Editor.

Channel 0 outputs are selected through a display control. Choices include an the image with green as alpha 0.0 (using the black swatch), colored backgrounds using the color swatch, or UV masks (display:1), alpha masks (display:2), or YUV (display 3,4,5 respectively) using the display selector.

If you can't afford a $400 muslim greenscreen, rumor has it some Blenderheads have even had fair results using a $4 tablecloth.

Iris Plugin Effect

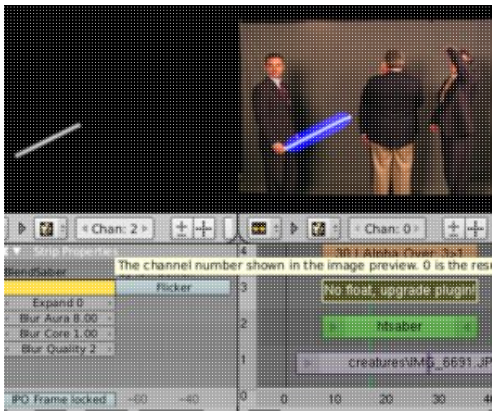


VSE Iris Plugin
This transition effect reveals the second strip while wiping the first by a decreasing circle, like the closing of an iris. Selecting In covers up the base image by revealing the second as the circle (iris) closes.

Jitter Plugin Effect

VSE Jitter Plugin
Too much caffeine? Can't hold the camera steady? Need to simulate Brownian motion? How about being time-phased out of sync by a harmonic distortion in the time-space continuum? Yeah, that sort of stuff happens to me every day too. But no one believes me, so use the Plugin to show them what it's like. Set the X and Y to be in pixels, and enable Random.

Lightsaber Plugin Effect



Lightsaber Animated Plugin
From: Intergalactic Indemnity Assurance Corporation

To: Claimant
Re: Request for Disability
Date: 15-05-3043 Terra-GMT

Dear Sir/Madam:

Our examiners have carefully reviewed the photographic evidence you submitted in connection with an unfortunate light-saber incident you claim happened during the filming of "Elephants Jedi" which resulted in the loss of your arm. After careful inspection, we have concluded that the so-called "lightsaber" was in fact a special effect which you created and staged to look like an accident waiting to happen. Your claim has been turned over to the authorities for prosecution.

Sincerely,

Jackaky Skinflint, CPA, AFP, EXT, CLIM, GAK
Claim Inspector #4395, Intergalactic Indemnity Assurance Corp.

WARNING: This message is confidential and for its intended recipient only. If you have received this message in error, please report to your local health office and submit to a free brain wipe. Thank you.

This effect does an edge detect on the base image, and adds a glow to the edge. In the example above, the scene has just a tube. The plugin makes that tube into a lightsaber, and the effect strip is AlphaOvered the live-action footage. Animating the tube to track his hand would make it look like he's about to lop my arm off. At least, that was the idea. Too bad Jackaky wasn't fooled.

LSD Plugin Effect



VSE LSD Plugin
Now you can experience the joy of mind-altering drugs without actually incurring any risk or harm, and saving a lot of money in the process. This self-animating plugin does bizzarre things to the colors of your image. There is not control over this effect; it does what it wants to (just like the real LSD). But, no flashbacks (unless you duplicate the source strip and effect).

OldMovie Plugin Effect



VSE Old Movie Plugin

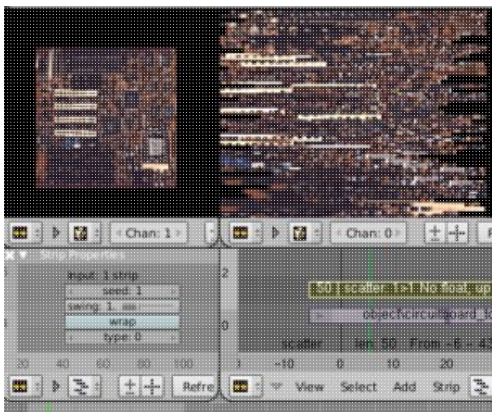
Back in the old dayz (I just love saying that), of black and white movies, the film stock was grainy, and the developer brushes would clog and scratch the negatives, resulting in long black streaks on the print. Real crap for quality. If you long for the old days, or want to post some Lost tapes from the original Dharma Initiative, use this effect to make them seem old. Thank you, and Good Luck. Oh, and if you do a good job, everyone will want to watch it again.

Red-Green Stereo Plugin Effect



VSE Red-Green Stereoscopic Viewer Plugin
Don't throw those Trix glasses away! Use two images, one taken from the persective of the left eye, another from the right, as the two images to feed this plugin. The Plugin generates a red-green stereoscopic image that appears 3-D when you put the glasses on. If you don't have a pair, you might be able to borrow indefinately (steal) your little sister's from one of her DVDs, or find a site on the web that gives them away.

Scatter Plugin Effect



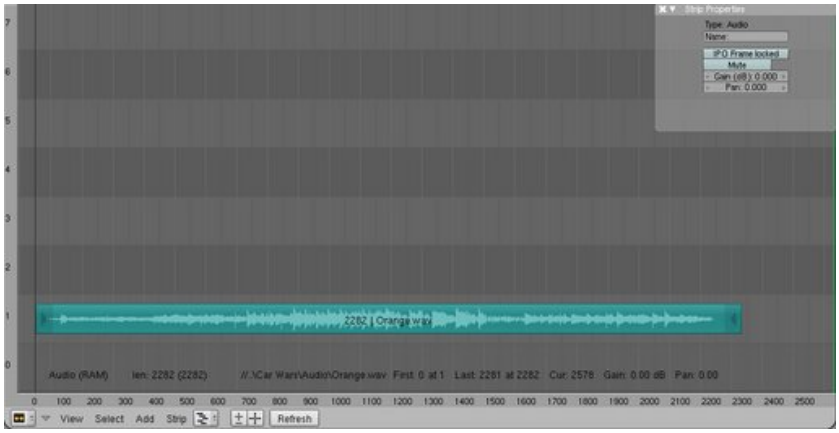
VSE Scatter Animated Plugin
So there we were, outbound from AlphaCentauri, when our Falx Confibulator board blew. We pulled the circuit board, and our logic tracer was showing us a problem somewhere in the VSE circuitry. Drifting just outside the nebulae, way too close to *the Fringe* for any of us to feel comfortable, we tried desparately to get Tech on-line to give us some assistance. "Hello, Tech? We have a problem here. Can you see the board?" we asked. "Barely" came the wiz's resonse, "I've got another one here and I'll walk you through..zzzcchdhitchchch" and the signal got garbled and full of static all of a sudden. "Must be an ion storm inside the nebula" my comms officer surmised. "Dammit, we need a clear signal! I can't tell what chip he's pointing to, and unless we get this Confibulator back on-line, we're stranded out here!" I yelled. With that, we

all turned to the viewport, and watched in quiet desparation as we drifted ever closer to *the Fringe*.

Warp Plugin Effect

First selecting a vector map (a black and white image that tells the plugin what to do) and an image to manipulate, the Warp plugin shifts the pixels sideways and vertically. Set the amount of the warp using the slider, and enable animate for a smooth warping over the distance of the video. Alternatively, use an Ipo curve to control the amount of the warp.

Sound Sequence Editor



A sound strip in the sequence editor.
Blender contains a multi-track Audio sequencing toolbox. You can add WAV files from your hard disk as a file, or as encoded within a movie, and mix them using an IPO Curve as a volument control.

Options

Mode: Video Sequence Editor

Hotkey: â § ShiftA

Menu: Add->Audio (RAM) or Add->Audio HD

Audio-RAM loads a file into memory and plays it from there. You can only load stand-alone WAV files. Audio-HD plays the sound back from the hard disk and thus does not take up memory. With Audio HD, you can load stand-alone WAV files, but also audio tracks from movies.

For either, a green audio strip will be created. With Audio RAM, a waveform is created that shows you the waveform inside the green strip, scaled to the height of the green strip. Since Audio-RAM files are read into memory, changing the audio file will not affect playback, and you will have to re-open the file so that Blender re-reads the file.

Hiss, Crackle and Pop
Some audiophile users report that Hiss is introduced sometimes if Audio RAM is used. There must be some decoding or sampling going on, that does not occur when Audio HD is used, that introduces some playback noise. If you hear pops and crackles, usually that is a sign that your hardware cannot keep up in real-time playback. They will not be present in your final rendered animation output (but they may show up in Game mode.

Audio Mixing in the VSE

You can have as many Audio strips as you wish and the result will be the mixing of all of them. You can give each strip its own name and Gain (in dB) via the N menu. This also let you set a strip to mute or 'Pan' it; -1 is hard left, +1 is hard right, with percentages in-between.

Overlapping strips are automatically mixed down during ANIM processing. For example, you can have the announcer on channel 5, background music on channel 6, and foley sound effects on channel 7.

Working with Audio Tracks

An audio track (strip) is just like any other strip in the VSE. You can grab and move it, adjust its starting offset using RMB over the arrow end handles, and K cut it into pieces. A useful example is cutting out the "um's" and dead voice time.

Animating Audio Track Properties

A 'Volume' IPO curve can be added for the selected audio strip by CtrlLMB clicking in the IPO Window as it is done for effect strips. The Fac channel controls the volume. You want to set a value somewhere between 0.0 and 1.0, and the volume becomes that percent; 0.6 is 60%. You can add a gain to the volume through the strip properties (N). You can make a curve by having multiple points, to vary the volume over its length. Press â Tab to edit the curve, just like any old bezier IPO curve.

IPO frames 1-100 correspond to the whole sample length. In the Y direction, 1.0 is full volume, 0.0 is completely silent. To make IPO frames match video frames, switch on "IPO frame locking" in the N-keys dialog. Only the FFMPEG-output system is currently able to mix audio and video into one output stream. Use CtrlLMB to add IPO control points, and â Tab to edit an IPO curve.

Animating an audio strip affects the volume of the strip in the resulting composite. Use IPO animation on an audio strip to fade in/out background music or to adjust volume levels. Layered/crossed audio strips are added together; the lower channel does not override and cut out higher channels. This makes Blender an audio mixer. By adding audio tracks and using the IPO curves to adjust each tracks' sound level, you have an automated dynamic multi-track audio mixer!

The output is therefore a video file if the ANIM button in the Anim Panel of the Scene Context/Render Sub-context is used as described before. An audio file may be created via the MIXDOWN button in the Sequencer button of the Scene Context, Sound Sub-context. This WAV file contains the full audio sequence and is created in the same directory of the video file and with the same name but with a .WAV extension. You can mix Video and Audio later on with an external program or by adding it to, for example, an image sequence strip as described above.

The advantage of using Blender's sequence editor lies in the easier synchronization attainable by sequencing frames and sound in the same application.

To enable audio synchronisation after importing an audio track, select the Scene button (F10) in the buttons window then choose the Sound Block Button (small blue sine wave). In here you'll see the Sync and Scrub tools.

- Sync lets Blender drop image frames to keep up with realtime audio when you play an animation in the 3D window. This gives you a rough overview of the timing of your animation.
- Scrub allows you to drag your frame-marker or change frames in any window and it will play a clip of audio for that point in time.

Draging the frame-marker over a range of frames in the Action Editor will allow you to hear roughly where specific sounds occur so that you can key poses or shapes on this frame.

Extending Blender

Unlike many programs you may be familiar with, Blender is not monolithic and static. You can extend its functionalities without having to modify the source and recompile. There are two ways of extending Blender: Python scripting and Binary Plugins.

Python scripts

Outside of the Blender executable, there are literally hundreds of scripts written by many people:

- Several scripts are bundled with Blender.
- Not all scripts can be included within the Blender release for logistical or support reasons, but many of them work reliably and are very useful, so they have been catalogued and are available in this wiki in the Scripts Catalog and in the new Extensions tracker.

Plugins

There are many plugins to extend Blender in two major areas:

- Texture Plugins - Change the way a material appears.
- Video Sequences - Do something neat to a video stream.

The default installation of Blender includes some plugins, but you can find others at Kent Meinâ← s repository and in the new Extensions tracker.

Saving your own Extensions

Files location

In a **GNU/Linux** distribution, there are two folders called â← `scripts`â← and â← `plugins`â← under one of these folders:

- `/usr/share/blender/`, OR
- `/usr/lib/blender/`,

depending on your distribution.

In a **Windows** installation, there are also two folders called â← `scripts`â← and â← `plugins`â← , but the â← `scripts`â← folder is in a directory called â← `.blender`â← , which is located alongside the `plugins` folder. The location of those folders depends on the choice you make during installation:

Use application data directory
: In this case `scripts` and `plugins` will be in a place like:
`C:\Users\****\AppData\Roaming\Blender Foundation\Blender\`

Use installation directory

In this case â← `plugins\â←` and â← `.blender\scripts\â←` will be in the installation directory of Blender.

I have defined a %HOME% variable - please install here
In this case the folders will be at the defined â← `%HOME%â←` location.

You can setup a path for scripts and one for plugins in User preferences (User Preferences â File Paths).

Scripts

If you find a script useful for you, download it to your `scripts` directory and it will be there when you need it.

If it is not in the `scripts` folder, you have to load it using the text editor window and then File â Run Python Script.

Plugins

If a plugin is in the Blender `plugins` install directory (or where you choose via User Preferences), you will easily find it when you use Texture or Sequence Plugin options.

If not, you will have to find it on your hard disk.

In both case it will be dynamically loaded for you and begin working immediately.

Python Scripts

Blender has a very powerful yet often overlooked feature. It exhibits an internal fully fledged Python interpreter. This allows any user to add functionalities by writing a Python script.

Python is an interpreted, interactive, object-oriented programming language. It incorporates modules, exceptions, dynamic typing, very high level dynamic data types, and classes. Python combines remarkable power with very clear syntax. It was expressly designed to be usable as an extension language for applications that need a programmable interface, and this is why Blender uses it.

Of the two main ways of extending Blender, the other one being binary plugins, Python scripting is more powerful, versatile yet easier to comprehend and robust. It is generally preferred to use Python scripting than writing a plugin.

Actually Python scripting had somewhat limited functionalities up to Blender 2.25, the last of NaN releases. When Open Sourcing Blender many of the new developers gathered around the Foundation elected to work on it and, together with UI change, Python API is probably the single part of Blender which got the greatest development. A full reorganization of what existed was carried out and many new modules added.

This evolution is still ongoing and even better integration is expected in forthcoming Blender versions.

- Setting up Python
- The Text Editor
- How to use scripts
- A working example
- Bundled Scripts
- References

Python Scripting

Please refer to the Python Development index

Downloading Python

To have access to the standard Python modules you need a complete working Python install. You can download this from the Python web site.

Setting the PYTHONPATH environment variable

Be sure to check on Blender web site which is the *exact* Python version which was built into Blender to prevent compatibility issues. Blender must also be made aware of *where* this full Python installation is. This is done by defining a `PYTHONPATH` environment variable.

Win95,98,Me

Once you have installed Python in, say, `C:\PYTHON22`, you must open the file `C:\AUTOEXEC.BAT` with your favourite text editor, add a line:

```
SET PYTHONPATH=C:\PYTHON22;C:\PYTHON22\DLLS;C:\PYTHON22\LIB;C:\PYTHON22\LIB\LIB-TK
```

and reboot the system.

WinNT,2000,XP

Once you have installed Python in, say, `C:\PYTHON22`, go on the â← My Computerâ← icon on the desktop, click RMB and select Properties. Select the Advanced tab and press the Environment Variables button.

Below the System Variables box (the second one), hit New. If you are not an administrator you might be unable to do that. In this case hit New in the upper box.

Now, in the Variable Name box, type `PYTHONPATH`, and in the Variable Value box,type:

```
C:\PYTHON22;C:\PYTHON22\DLLS;C:\PYTHON22\LIB;C:\PYTHON22\LIB\LIB-TK
```

Hit OK repeatedly to exit from all dialogs. You may or may not have to reboot, depending on the OS.

Linux and other UNIXes

Normally you will have Python already there. If not, install it. You will have to discover where it is. This is easy, just start a Python interactive shell by opening a shell and by typing `python` in there. Type the following commands:

```
>>> import sys
 >>> print sys.path
```

and note down the output, it should look like:

```
['', '/usr/local/lib/python2.2', '/usr/local/lib/python2.2/plat-linux2',
 '/usr/local/lib/python2.0/lib-tk', '/usr/local/lib/python2.0/lib-dynload',
 '/usr/local/lib/python2.0/site-packages']
```

Add this to your favorite `rc` file as an environment variable setting. For example, add in your `.bashrc` the line:

```
export PYTHONPATH=/usr/local/lib/python2.2:/usr/local/lib/python2.2/plat-linux2:
                  /usr/local/lib/python2.2/lib-tk:/usr/local/lib/python2.2/lib-dynload:
                  /usr/local/lib/python2.0/site-packages
```

(**warning**: all on a *single line*). Open a new login shell, or logout and login again.

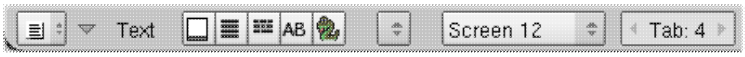
Note
Even though the `PYTHONPATH` environment variable is not defined, Blender seems to have no problem finding Python, when both have been installed with the package system of the distribution. Probably because Python is in a standard place (e.g. `/usr/lib/python2.5`).

The Text Editor

Blender has a Text Editor among its windows types, accessible via the Text Editor button (

) of the Window type menu, or via â § ShiftF11.

The newly opened Text window is grey and empty, with a very simple toolbar (*Text Toolbar*). From left to right there are the standard Window type selection button and the window menus. Then, five toggle buttons: the full screen switch, one which shows/hides the line numbers for the text, one which enables/disables the word wrapping, one which enables/disables the syntax highlighting, and the last one which enables/disables the Python text plugins. At the end, a dropdown menu to select/add/open a text buffer, and two settings: the font of the editor, and the number of spaces in a tabulation.



Text Toolbar.

The dropdown menu () allows you to select which Text buffer is to be displayed, as well as allowing you to create a new buffer or load a text file. If you choose to load a file the Text Editor temporarily becomes a File Browser window, with the usual functions. Once a text buffer is in the Text window, this behaves as a quite simple text editor.

Typing on the keyboard produces text in the text buffer. As usual, pressing dragging and releasing LMB selects text. The following keyboard commands apply:

- AltC or CtrlC - Copies the marked text into the text clipboard.
- AltX or CtrlX - Cuts out the marked text into the text clipboard.
- AltV or CtrlV - Pastes the text from the clipboard to the cursor in the Text window.

Blenderâ ← s cut/copy/paste clipboard is *separate* from the operating system (OS) clipboard. So normally you *cannot* cut/paste/copy out from/into Blender. To access your OS clipboard:

- Ctrlâ § ShiftC - To copy text to the OS buffer (e.g. if you want to paste that text into another application).
- Ctrlâ § ShiftX - To cut and copy text to the OS buffer.
- Ctrlâ § ShiftV - To paste text from the OS buffer (e.g. you copied some text from your web browser or document editor).
- AltS - Saves the text as a text file, a File Browser window appears.
- AltO - Loads a text, a File Browser window appears.
- AltF - Pops up the Find & Replace toolbox.
- â § ShiftAltF or RMB - Pops up the Text menu for the Text window.
- AltJ - Pops up a Num control where you can specify a line number the cursor will jump to.
- AltP - Executes the text as a Python script.
- AltU - Undo.
- AltR - Redo.
- CtrlR - Reopen (reloads) the current buffer (all non-saved modifications are lost).
- AltM - Converts the content of the text window into 3D text (max 100 chars).

To delete a text buffer just press the X button next to the bufferâ ← s name, just as you do for materials, etc. The most notable keystroke is AltP which makes the content of the buffer being parsed by the internal Python interpreter built into Blender. The next page will present an example of Python scripting. Before going on it is worth noticing that Blender comes with only the bare Python interpreter built in, and with a few Blender-specific modules, those described in the API references.

The Text Editor has now also some dedicated Python scripts, which add some useful writing tools, like a class/function/variable browser, completionâ←¦ You can access them through the Text â Text Plugins menu entry.

Other usages for the Text window

The text window is handy also when you want to share your .blend files with the community or with your friends. A Text window can be used to write in a README text explaining the contents of your blender file. Much more handy than having it on a separate application. Be sure to keep it visible when saving! If you are sharing the file with the community and you want to share it under some license you can write the license in a text window.

A working Python example

Now that youâ← ve seen that Blender is extensible via Python scripting and that youâ← ve got the basics of script handling and how to run a script, before smashing your brain with the full python API reference, letâ← s have a look at a quick working example.

We will present a tiny script to produce polygons. This indeed duplicates somewhat the Space â Add â Mesh â Circle toolbox option, but will create â← filledâ← polygons, not just the outline.

To make the script simple yet complete it will exhibit a Graphical User Interface (GUI) completely written via Blenderâ← s API.

Headers, importing modules and globals

The first 32 lines of code are listed below:

Script header

```
1. ##########################################################
2. #
3. # Demo Script for Blender 2.3 Guide
4. #
5. ##########################################################S68
6. # This script generates polygons. It is quite useless
7. # since you can do polygons with ADD->Mesh->Circle
8. # but it is a nice complete script example, and the
9. # polygons are 'filled'
10. ##########################################################
11.
12. ##########################################################
13. # Importing modules
14. ##########################################################
15.
16. import Blender
17. from Blender import NMesh
18. from Blender.BGL import *
19. from Blender.Draw import *
20.
21. import math
22. from math import *
23.
24. # Polygon Parameters
25. T_NumberOfSides = Create(3)
26. T_Radius = Create(1.0)
27.
28. # Events
29. EVENT_NOEVENT = 1
30. EVENT_DRAW = 2
31. EVENT_EXIT = 3
```

After the necessary comments with the description of what the script does there is (lines `[016-022]`) the importing of Python modules. `Blender` is the main Blender Python API module. `NMesh` is the module providing access to Blenderâ← s meshes, while `BGL` and `Draw` give access to the OpenGL constants and functions and to Blenderâ← s windowing interface, respectively.

The `math` module is Pythonâ← s mathematical module, but since both the `math` and the `os` modules are built into Blender you donâ← t need a full Python install for this!

The polygons are defined via the number of sides they have and their radius. These parameters have values which must be defined by the user via the GUI, hence lines [025-026] create two â← generic buttonâ← objects, with their default starting value.

Finally, the GUI objects works with, and generates, events. Events identifier (or events ID for short) are integers left to the coder to define (they should all be different, unless special needs!). It is usually a good practice to define mnemonic names for events, as is done here in lines [029-031].

Drawing the GUI

The code responsible for drawing the GUI should reside in a `draw` function (*GUI drawing*).

GUI drawing

```
33. ########################################################
34. # GUI drawing
35. ########################################################
36. def draw():
37. global T_NumberOfSides
38. global T_Radius
39. global EVENT_NOEVENT,EVENT_DRAW,EVENT_EXIT
40. 
41. ########## Titles
42.         glClear(GL_COLOR_BUFFER_BIT)
43.         glRasterPos2d(8, 103)
44.         Text("Demo Polygon Script")
45. 
46. ######### Parameters GUI Buttons
47.         glRasterPos2d(8, 83)
48.         Text("Parameters:")
49.         T_NumberOfSides = Number("No. of sides: ", EVENT_NOEVENT, 10, 55, 210, 18,
50.                 T_NumberOfSides.val, 3, 20, "Number of sides of out polygon");
51.         T_Radius = Slider("Radius: ", EVENT_NOEVENT, 10, 35, 210, 18,
52.                 T_Radius.val, 0.001, 20.0, 1, "Radius of the polygon");
53. 
54. ######### Draw and Exit Buttons
55.         Button("Draw",EVENT_DRAW , 10, 10, 80, 18)
56.         Button("Exit",EVENT_EXIT , 140, 10, 80, 18)
```

Lines [037-039] merely grant access to global data. The real interesting stuff starts from lines [042-044]. The OpenGL window is initialized, and the current position set to (`x=8, y=103`). The origin of this reference is the lower left corner of the script window. Then the title â← `Demo Polygon Script`â← is printed. A further string is written (lines [047-048]).

Then the input buttons for the parameters are created. The first (lines [049-050]) is a Num Button, exactly like those in the various Blender Button Windows. For the meaning of all the parameters please refer to the API reference. Basically there is the button label, the event generated by the button, its location (x,y) and its dimensions (width, height), its value, which is a data belonging to the Button object itself, the minimum and maximum allowable values and a text string which will appear as a help while hovering on the button, as a tooltip. Lines [051-052] defines a Num Button with a slider, with a very similar syntax. Lines [055-056] finally create a `Draw` button which will create the polygon and an `Exit` button.

Managing Events.

The GUI is not drawn, and will not work, until a proper event handler is written and registered (*Handling events*).

Handling events

```
def event(evt, val):
if (evt == QKEY and not val):
                Exit()

def bevent(evt):
global T_NumberOfSides
global T_Radius
global EVENT_NOEVENT,EVENT_DRAW,EVENT_EXIT

######### Manages GUI events
if (evt == EVENT_EXIT):
                Exit()
elif (evt== EVENT_DRAW):
                Polygon(T_NumberOfSides.val, T_Radius.val)
                Blender.Redraw()

Register(draw, event, bevent)
```

Lines `[058-060]` define the keyboard event handler, here responding to the Q with a plain `Exit()` call.

More interesting are lines `[062-072]`, in charge of managing the GUI events. Every time a GUI button is used this function is called, with the event number defined within the button as a parameter. The core of this function is hence a â← selectâ← (if/elif) structure executing different codes according to the event number.

As a last call, the `Register()` function is invoked. This effectively draws the GUI and starts the event capturing cycle (â← event loopâ←).

Mesh handling

Finally, *Main function* shows the main function, the one creating the polygon. It is a rather simple mesh editing, but shows many important points of the Blenderâ← s internal data structure.

Main function

```
##########################################################
# Main Body
##########################################################
def Polygon(NumberOfSides,Radius):

######### Creates a new mesh
        poly = NMesh.GetRaw()

######### Populates it of vertices
for i in range(0,NumberOfSides):
                phi = 3.141592653589 * 2 * i / NumberOfSides
                x = Radius * cos(phi)
                y = Radius * sin(phi)
                z = 0

                v = NMesh.Vert(x,y,z)
                poly.verts.append(v)

######### Adds a new vertex to the center
        v = NMesh.Vert(0.,0.,0.)
        poly.verts.append(v)

######### Connects the vertices to form faces
for i in range(0,NumberOfSides):
```

```
100.                f = NMesh.Face()
101.                f.v.append(poly.verts[i])
102.                f.v.append(poly.verts[(i+1)%NumberOfSides])
103.                f.v.append(poly.verts[NumberOfSides])
104.                poly.faces.append(f)
105.
106. ######### Creates a new Object with the new Mesh
107.        polyObj = NMesh.PutRaw(poly)
108.
109.        Blender.Redraw()
```

The first important line here is number `[082]`. Here a new mesh object, `poly` is created. The mesh object is constituted of a list of vertices and a list of faces, plus some other interesting stuff. For our purposes the vertices and faces lists are what we need.

Of course the newly created mesh is empty. The first cycle (lines `[085-092]`) computes the x,y,z location of the `NumberOfSides` vertices needed to define the polygon. Being a flat figure it is z=0 for all. Line `[091]` calls the `NMesh` method `Vert()` to create a new vertex object of co-ordinates (x,y,z). Such an object is then appended (line `[092]`) in the `poly` Mesh `verts` list.

Finally (lines `[095-096]`) a last vertex is added in the centre.

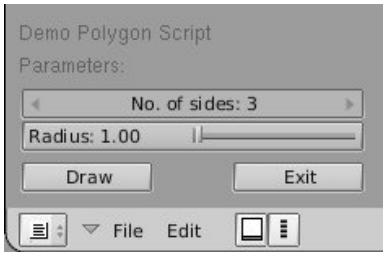
Lines `[099-104]` now connect these vertices to make faces. It is not required to create all vertices beforehand and then faces. You can safely create a new face as soon as all its vertices are there.

Line `[100]` creates a new face object. A face object has its own list of vertices `v` (up to 4) defining it.

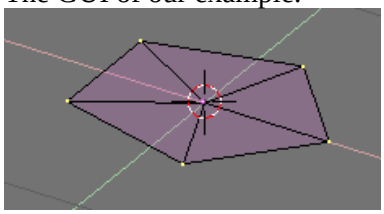
Lines `[101-103]` appends three vertices to the originally empty `f.v` list. The vertices are two subsequent vertices of the polygon and the central vertex. These vertices must be taken from the Mesh `verts` list.

Finally line `[104]` appends the newly created face to the `faces` list of our `poly` mesh.

Conclusions



The GUI of our example.



The result of our example script.

If you create a `polygon.py` file containing the above described code and load it into a Blender text window, as you learned in the previous page, and press AltP in that window to run it, you will see the script disappearing and the window turn grey. In the lower left corner the GUI will be drawn (*The GUI of our example*).

By selecting, for example, 5 vertices and a radius 0.5, and by pressing the Draw button a pentagon will appear on the xy plane of the 3D window (*The result of our example script*).

Blenderâ← s Python API

The full Python API (Application Programmer Interface) of Blender is documented here:

2.42 API 2.43 API 2.44 API 2.45 API 2.46 API 2.47 API 2.48 API 2.49 API

Scripts

There are more than one hundred different scripts for Blender available on the net.

As with plugins, scripts are very dynamic, changing interface, functionalities and web location fairly quickly, so for an updated list and for a live link to them please refer to one of the two main Blender sites:

- www.blender.org
- www.blenderartists.org
- Python extensions on this wiki.

There are many scripts that are packaged with Blender, as well as hundreds available for downloading from individual sites. To facilitate good documentation of these resources, an entirely new wiki section has been established. Please refer to that page for more information about bundled and available scripts that extend the functionality of Blender.

Python scripts can, after a long and stable history, be incorporated into the Blender menu system and appear as options on a windowâ← s menu option. This bundling makes them readily accessible to the user. Here are some windows that have/use scripts:

- User Preferences
- 3D View
- UV/Image Editor
- Text Editor
- Node Editor

This section provides specific links to further documentation on the specific Python Scripts that are available from Blender menus. For more information on all known scripts, see the catalog.

Bundled Scripts

User Preferences Window

File â Import

Youâ← ll find more doc about importers here.
This is a list of support for importers within Blender (Note : it definitively needs an update, there are much more import scripts, nowâ←¦).
NA (not applicable), unsupported by the format.
YES/NO or ***** for notes.

Native	Ext	Object	Mesh	Mat	UV	Vert Color	Arm	Animation	Skin	Scene	Lgt	Cam	World
3D Studio	3DS	YES	YES	YES	YES	NO	NO	NO	YES	NA	NO	NO	NA
AC3D	AC	YES	YES	YES	YES	NA	NA	NA	NA	NA	NA	NA	NA
Quake 2	MD2	YES	YES	NA	YES	NA	NA	YES	YES	NA	NA	NA	NA
Wavefront OBJ	OBJ	YES	YES	YES	YES	NA	NA	NA	NA	NA	NA	NA	NA
Native	Ext	Object	Mesh	Mat	UV	Vert Color	Arm	Animation	Skin	Scene	Lgt	Cam	World

File â Export

Youâ← ll find more doc about exporters here.
This is a list of support for exporters within Blender (Note : it definitively needs an update, there are much more export scripts, nowâ←¦).
NA (not applicable), unsupported by the format.
YES/NO or ***** for notes.

Native	Ext	Object	Mesh	Mat	UV	Vert Color	Arm	Animation	Skin	Scene	Lgt	Cam	World
3D Studio	3DS	YES	YES	YES	YES	NO	NO	NO	YES	NA	NO	NO	NA
AC3D	AC	YES	YES	YES	YES	NA	NA	NA	NA	NA	NA	NA	NA
Quake 2	MD2	YES	YES	NA	YES	NA	NA	YES	YES	NA	NA	NA	NA
Wavefront OBJ	OBJ	YES	YES	YES	YES	NA	NA	NA	NA	NA	NA	NA	NA
Native	Ext	Object	Mesh	Mat	UV	Vert Color	Arm	Animation	Skin	Scene	Lgt	Cam	World

Render

- Backface Cull.

Help

- Blender/Python Scripting API - Opens HTML browser to an internet page on Python in Blender.
- Getting Started - Opens HTML browser to an internet page for newbies.
- HotKey and MouseAction Reference - Changes a window to a quick reference.
- ID Property Browser - Sets/inspects properties on any Blender object.
- Manual - Opens HTML browser to this wiki.
- Release Notes - Whatâ← s new in this Blender version.
- Scripts Help Browser - Gives you some info and help on scripts themselves.

- Tutorials - Points you to tutorials that give step-by-step simple instructions on how to do great things with Blender.

Help â Websites

These scripts simply open an HTML page in your HTML browser (Firefox, IE, â←¦) and display the relevant page at www.blender.org:

- E-Shop - Get your Blenderware here!
- Blender Website.
- Developer Community - Programmers developing Blender.
- User Community - Users helping other users.

Help â System

- System Information... - A small info tool about your Blender, Python and OpenGL installations.

3D View Window

The available scripts that are shown on the menu vary by the mode you are in. For example, the list of scripts that are usable when editing a mesh in Edit mode are different than the scripts available when working with whole Objects.

Object mode â Object â Scripts

- Apply Deformation.
- Axis Orientation Copy.
- Batch Object Name Edit.
- Bone Weight Copy.
- Cookie Cut from View.
- Data Copier.
- Drop Onto Ground.
- Find by Data Use.
- Object Name Editor.
- Randomize Loc Size Rot.
- Selection to DupliGroup.
- UV Copy from Active.

Edit mode â Mesh â Scripts

- Bevel Center.
- Clean Meshes.
- Discombobulator.
- Edges to Curve.
- Mirror Vertex Locations & Weight.
- Poly Reducer.
- Rvk1 to Rvk2.
- Solid Wireframe.
- Solidify Selection.
- Unfold.
- Unweld vertex/ices.

Vertex Paint mode â Paint

- VCol Gradient.
- Self Shadow VCols (AO).
- Copy from Material.

Weight Paint mode â Paint

- Clean Weight.
- Copy Active Group.
- Envelope via Group Objects.
- Normalize/Scale Weight.
- Grow/Shrink Weight.
- Weight Gradient.

UV/Image Editor Window

Image mode

These scripts help you when working with images that are probably going to be used as UV-mapped images.

- Fix Broken Paths.
- Edit Externally.
- Consolidate into one image.
- Billboard Render on Active.

UVs mode

- Save UV Face Layout.
- Seams from Islands.

Text Editor Window

These scripts helps you when writing a Python script! Youâ← ll find more doc about them here.

Text â Text Plugins
â←¦

Text â Script Templates
â←¦

Making a Script â← Bundledâ←

To make a script registered within blender (to show up in a menu) you need to include a script header. Please see â← Registering Scriptsâ← on this page Python API Page.

Additional Resources

Please also see the BlenderArtists Python forum.

Blenderâ← s Plugins System

This section reports an in-depth reference for Blenderâ← s Texture and Sequence plugins:

Texture plugins

- Texture plugins list (from *Texture* section)
- Texture plugins specifications

Sequence plugins

- Sequencer plugins list (from *Sequencer* section)
- Sequence plugins specifications

Writing a Texture Plugin

In this Section we will write a basic texture plugin and then go through the steps to use a texture plugin.

The basics behind a texture plugin is that you are given some inputs; position, and normal values as well as some other info. Then you return intensity, colour and/or normal information depending on the type of texture plugin.

All the files necessary to develop plugins as well as a few sample plugins can be found in the `blender/plugins` directory. You can alternately get a bunch of plugins from this plugins repository. Plugins are supported (loaded/called) in Blender using the `dlopen()` family of calls. For those unfamiliar with this system, it allows a program (Blender) to use a compiled object as if it were part of the program itself, similar to dynamically linked libraries, except the objects to load are determined at runtime.

The advantage of using the `dlopen` system for plugins is that it is very fast to access a function, and there is no overhead in interfacing to the plugin, which is critical when as (in the case of texture plugins) the plugin can be called several million times in a single render. The disadvantage of the system is that the plugin code works just like it is part of Blender itself, if the plugin crashes, Blender crashes.

The include files found in the `plugin/include/` subdirectory of the Blender installation document the Blender functionality provided to the plugins. This includes the `Imbuf` library functions for loading and working with images and image buffers, and noise and turbulence functions for consistent texturing.

Specification

```
#include <plugin.h>
```

Every Blender plugin should include this header file, which contains all of the structures and defines needed to properly work with Blender.

```
char name[]="Tiles";
```

A character string containing the plugin name, this value will be displayed for the textureâ← s title in the Buttons window, Texture sub-context.

```
#define NR_TYPES 2; char stnames[NR_TYPES][16]= {"Square", "Deformed"};
```

Plugins are allowed to have separate subtypes for minor variations on algorithms - for example the default clouds texture in Blender has the â← Defaultâ← and â← Colorâ← subtypes. `NR_TYPES` should be defined to the number of subtypes required by your plugin, and a name for each subtype should be given (`stnames[NR_TYPES]`). Every plugin should have at least 1 subtype and a subtype name.

```
VarStruct varstr[]= {...};
```

The `varstr` contains all of the information Blender needs to display buttons for a plugin. Buttons for plugins can be numerical for input data, or text for comments and other information. Plugins are limited to a maximum of 32 variables. Each `VarStruct` entry consists of a type, name, range information, and a tool tip.

The **type** defines the data type for each button entry, and the way to display the button. For number buttons this value should be a combination (binary OR, â← `|`â←) of â← `INT`â← or â← `FLO`â← for the number format (integer or float), and â← `NUM`â← , â← `NUMSLI`â← , or

â← TOGâ← , for the button type (numeric, numeric with slider, or toggle). Text buttons should have a type of â← LABELâ← .
The **name** is what will be displayed on (or beside) the button. This is limited to 15 characters.
The **range information** consists of three floats that define the default, minimum, and maximum values for the button. For TOGgle buttons the minimum is set in the pressed state, and the maximum is set in the depressed state.
The **tool tip** is a string that will be displayed when the mouse is over this button (if the user has Tool Tips on). This has a limit of 80 characters, and should be set to the NULL string ("") if unused.

```
typedef struct Cast {...};
```

The Cast structure is used in calling the doit function, and serves as a way to simply access to each pluginâ← s data values. The cast should contain, in order, an integer or float for every button defined in the varstr, including text buttons. Typically these should have the same name as the button for simple reference.

```
float result[8];
```

The result array is used to pass information to and receive information from the plugin. The result values are mapped as follows:

result Index	Significance	Range
result[0]	Intensity value	0.0 to 1.0
result[1]	Red color value	0.0 to 1.0
result[2]	Green color value	0.0 to 1.0
result[3]	Blue color value	0.0 to 1.0
result[4]	Alpha color value	0.0 to 1.0
result[5]	X normal displacement	-1.0 to 1.0
result[6]	Y normal displacement	-1.0 to 1.0
result[7]	Z normal displacement	-1.0 to 1.0

The plugin should always return an intensity value. Returning RGB or a normal are optional, and should be indicated by the doit() return flag â← 1â← (RGB) or â← 2â← (Normal).
Before the plugin is called, Blender includes the current rendering-normal in result[5], result[6] and result[7].

`float cfra`
The cfra value is set by Blender to the current frame before every render pass. This value is the frame number (Â± 0.5 depending on the field settings).

`plugin_tex_doit` prototype
: The `plugin_tex_doit` function should be prototyped for use by the `getinfo` function. You do not need to change this line.

`plugin_tex_getversion`
: This function must be in each plugin for it to be loaded correctly. You should not change this function.

`plugin_but_changed`
: This function is used to pass information about what buttons the user changes in the interface. Most plugins should not need to use this function, only when the interface allows the user to alter some variable that forces the plugin to do recalculation (a random hash table for example).

`plugin_init`
: If needed plugins may use this function to initialize internal data.
Note: This init function can be called multiple times if the same plugin texture is copied. Do not init global data specific to a single instance of a plugin in this function.

`plugin_getinfo`
: This function is used to communicate information to Blender. You should never need to change it.

`plugin_tex_doit`
: The doit function is responsible for returning information about the requested pixel to Blender.
The Arguments
`int stype`
: This is the number of the selected subtype, see the `NR_TYPES` and `char stnames[NR_TYPES][16]` entries above.
`Cast *cast`
: The `Cast` structure which contains the plugin data, see the `typedef struct Cast` entry above.
`float *texvec`
: This is a pointer to three floats, which are the texture coordinates for which a texture value is to be returned.
`float *dxt, float *dyt`
: If these pointers are non-NULL they point to two vectors (two arrays of three floats) that define the size of the requested texture value in pixel space. They are only non-NULL when OSA is on, and are used to calculate proper anti aliasing.
The `doit` function should fill in the result array (it should *always* fill in an intensity value), and return 0, 1, 2 or 3 depending on what values have been filled in:
- **0** for intensity only,
- **1** for intensity + RGBA colour,
- **2** for intensity + normal, and
- **3** for intensity + RGBA colour + normal.

Texture/Material Interaction

Blender is somewhat different from most 3D packages in the logical separation between textures and materials. In Blender textures are objects that return certain values, signal generators in fact (like wave generators in synthesizer, for example). Materials control the mapping of textures onto objects, what is affected, how much, in what way, etc.

Properly designed plugins should only include variables to affect the signal returned, not the mapping of it. Buttons to control scale, range, axis, etc. are best only included when they make the texture easier to use (in the case of the Size button in the Tiles plugin), or when they speed up the calculation (the Intensity/Color/Bump subtypes in the Clouds2 plugin). Otherwise the Material sub-context make these buttons redundant, and the interface becomes needlessly complex.

Generic Texture Plugin

```
#include "plugin.h"

/* Texture name */
char name[24]= "";

#define NR_TYPES 3
char stnames[NR_TYPES][16]= {"Intens", "Color", "Bump"};

/* Structure for buttons,
 * {butcode, name, default, min, max, tooltip}
 */
VarStruct varstr[]= {
{NUM|FLO, "Const 1", 1.7,    -1.0, 1.0, ""},
};

typedef struct Cast {
float a;
} Cast;

float result[8];
float cfra;

int plugin_tex_doit(int, Cast*, float*, float*, float*);

/* Fixed Functions */
int plugin_tex_getversion(void) {
return B_PLUGIN_VERSION;
}

void plugin_but_changed(int but) { }
void plugin_init(void) { }

void plugin_getinfo(PluginInfo *info) {
    info->name= name;
    info->stypes= NR_TYPES;
    info->nvars= sizeof(varstr)/sizeof(VarStruct);
    info->snames= stnames[0];
    info->result= result;
    info->cfra= &cfra;
    info->vars= varstr;
    info->init= plugin_init;
    info->texdoit= (TexDoit) plugin_tex_doit;
    info->callback= plugin_but_changed;
}

int plugin_tex_doit(int stype, Cast *cast, float *texvec, float *dxt, float *dyt) {
if (stype == 1) {
return 1;
}
if (stype == 2) {
return 2;
}
return 0;
}
```

Our Modifications

The first step is to come up with a game plan. What is this plugin going to do, how are the users going to interact with it. For this example we will create a simple texture that creates a simple brick/block pattern.

Now weâ← ll copy our generic plugin to `cube.c` and will fill in the gaps. Its always a (very) good idea to add some comments. First off tell users what the plugin does, where they can get a copy, who they should contact for bugs/improvements, and any licensing restrictions on the code.

When using comments make sure you use â← `/* */`â← comments (C style). The plugins are in C and some C compilers do not accept â← `//`â← comments (C++ style).

```
/*
Description: This plugin is a sample texture plugin that creates a simple
brick/block pattern with it.
It takes two values a brick size, and a mortar size.
The brick size is the size of each brick.
The mortar size is the mortar size in between bricks.
Author: Kent Mein (mein@cs.umn.edu)
Website: http://www.cs.umn.edu/~mein/blender/plugins
Licensing: Public Domain
Last Modified: Tue Oct 21 05:57:13 CDT 2003
*/
```

Next we need to fill in the `name`, you should really keep this the same as your `.c` file, preferably descriptive, less than 23 chars, no spaces, and all lowercase.

```
char name[24]= "cube.c";
```

We are going to keep this plugin simple, and only have one type that deals with intensity. So we need the following:

```
#define NR_TYPES 1
char stnames[NR_TYPES][16]= {"Default"};
```

For our user interface, we are going to allow people to change the size of the brick and mortar, as well as the intensity values returned for the brick and mortar. For that we need to edit `varstr` and `Cast`. `Cast` should have a variable for each entry in `varstr`.

```
/* Structure for buttons,
 *      {butcode, name,     default, min, max, tooltip]
 */
VarStruct varstr[]= {
{NUM|FLO, "Brick",      0.8, 0.1, 1.0, "Size of Cell"},
{NUM|FLO, "Mortar",     0.1, 0.0, 0.4, "Size of boarder in cell"},
{NUM|FLO, "Brick Int",  1.0, 0.0, 1.0, "Color of Brick"},
{NUM|FLO, "Mortar Int", 0.0, 0.0, 1.0, "Color of Mortar"},
};

typedef struct Cast {
float brick, mortar, bricki, mortari;
} Cast;
```

Now we need to fill in plugin_tex_doit, we basically want to break down our texture into â← cellsâ← which will consist of a brick and the mortar along the bottom edges of that brick. Then determine if we are in the brick or the mortar. The following code should do that:

```
int plugin_tex_doit(int stype, Cast *cast, float *texvec, float *dxt, float *dyt) {
int c[3];
float pos[3], cube;

/* setup the size of our cell */
     = cast->brick + cast->mortar;

/* we need to do is determine where we are inside of the current brick. */
 [0] = (int)(texvec[0] / cube);
 [1] = (int)(texvec[1] / cube);
 [2] = (int)(texvec[2] / cube);
   [0] = ABS(texvec[0] - (c[0] * cube));
   [1] = ABS(texvec[1] - (c[1] * cube));
   [2] = ABS(texvec[2] - (c[2] * cube));

/* Figure out if we are in a mortar position within the brick or not. */
if ((pos[0] <= cast->mortar) ||
    (pos[1] <= cast->mortar) ||
    (pos[2] <= cast->mortar)) {
      [0] = cast->mortari;
} else {
      [0] = cast->brick;
}
return 0;
}
```

One thing to note, the ABS function (absolute value) is defined in a header in `plugins/include`. There are some other common functions there as well, be sure to take a look at whatâ← s there.

Compiling

â← bmakeâ← is a simple utility (shell script) to aid in the compilation and development of plugins. It can be found in the `plugins/` sub-directory of the Blender installation directory. It is invoked by: `bmake (`*`plugin_name`*`.c)`, and will attempt to link the proper libraries and compile the specified C file properly for your system.

If you are trying to develop plugins on a windows machine, bmake may not work for you. In that case you should look into using â← lccâ← . You can use the following to compile a plugin with lcc:

Assuming you have your plugins in `c:\blender\plugins`, here is an example of how you would compile the texture plugin `sinus.c`. Open a DOS prompt and do the following (note: Youâ← ll want to make sure the `lcc\bin` directory is in your path):

```
cd c:\blender\plugins\texture\sinus
lcc -Ic:\blender\plugins\include sinus.c
lcclnk -DLL sinus.obj c:\blender\plugins\include\tex.def
implib sinus.dll
```

The latest version of lcc does not have â← implibâ← . Instead, do the following:

```
cd c:\blender\plugins\texture\sinus
lcc -Ic:\blender\plugins\include sinus.c
lcclnk -dll -nounderscores sinus.obj c:\blender\plugins\include\tex.def
```

Note that `tex.def` is not distributed with the Blender source but is available from the Blender plug-in repository. Alternatively, create a `tex.def` file and copy and paste it the following :

```
EXPORTS
LibMain@12
```

```
plugin_but_changed
plugin_getinfo
plugin_init
plugin_tex_doit
plugin_tex_getversion
```

Writing a Sequence Plugin

In this Section we will write a basic sequence plugin and then go through the steps to use a sequence plugin.

The basics behind a sequence plugin are you are given as inputs 1 to 3 image buffers, as well as some other information, and you output a resulting image buffer.

All the files necessary to develop plugins as well as a few sample plugins can be found in the `blender/plugins` directory. You can alternately get a bunch of plugins from here.

Specification

```
#include <plugin.h>
```

Every Blender plugin should include this header file, which contains all of the structures and defines needed to properly work with Blender.

```
char name[]="Blur";
```

A character string containing the plugin name.

```
VarStruct varstr[]= {...};
```

The varstr contains all of the information Blender needs to display buttons for a plugin. Buttons for plugins can be numerical for input data, or text for comments and other information. Plugins are limited to a maximum of 32 variables. Each VarStruct entry consists of a type, name, range information, and a tool tip.

`type`
: defines the data type for each button entry, and the way to display the button.
For number buttons this value should be a combination (binary OR, â← |â←) of â← `INT`â← or â← `FLO`â← for the number format (integer or float), and â← `NUM`â← , â← `NUMSLI`â← , or â← `TOG`â← , for the button type (numeric, numeric with slider, or toggle).
Text buttons should have a type of â← `LABEL`â← .

`name`
: is what will be displayed on (or beside) the button. This is limited to 15 characters.

`range information`
: consists of three floats that define the default, minimum, and maximum values for the button. For `TOG`gle buttons, the minimum is set in the pressed state, and the maximum is set in the depressed state.

`tooltip`
: is a string that will be displayed when the mouse is over this button (if the user has Tool Tips on). This has a limit of 80 characters, and should be set to the NULL string (`""`) if unused.

```
typedef struct Cast {...};
```

The `Cast` structure is used in calling the `doit` function, and serves as a way to simply access each pluginâ← s data values. The cast should contain, in order, an integer or float for every button defined in the varstr, including text buttons. Typically these should have the same name as the button for simple reference.

```
float cfra
```

The `cfra` value is set by Blender to the current frame before every render pass. This value is the frame number (Â± 0.5 depending on the field settings).

`plugin_seq_doit` prototype
: The `plugin_seq_doit` function should be prototyped for use by the `getinfo` function. You do not need to change this line.

`plugin_seq_getversion`
: This function must be in each plugin for it to be loaded correctly. You should not change this function.

`plugin_but_changed`
: This function is used to pass information about what buttons the user changes in the interface. Most plugins should not need to use this function, only when the interface allows the user to alter some variable that forces the plugin to do recalculation (a random hash table for example).

`plugin_init`
: If needed plugins may use this function to initialize internal data.
Note: This init function can be called multiple times if the same plugin sequence is copied. Do not init global data specific to a single instance of a plugin in this function.

`plugin_getinfo`
: This function is used to communicate information to Blender. You should never need to change it.

`plugin_seq_doit`
: The sequence `doit` function is responsible for applying the pluginâ← s effect and copying the final data into the out buffer.
The Arguments
`Cast *cast`
: The `Cast` structure which contains the plugin data, see the `typedef struct Cast` entry above.
`float facf0`
: The value of the pluginâ← s Ipo curve for the first field offset. If the user hasnâ← t made an Ipo curve, this ranges between 0 and 1 for the duration of the plugin.
`float facf1`
: The value of the pluginâ← s Ipo curve for the second field offset. If the user hasnâ← t made an Ipo curve, this ranges between 0 and 1 for the duration of the plugin.
`int x, int y`
: The width and height of the image buffers, respectively.
`Imbuf *ibuf1`
: A pointer to the first image buffer the plugin is linked to. This will always be a valid image buffer.
`Imbuf *ibuf2`
: A pointer to the second image buffer the plugin is linked to. Plugins using this buffer should check for a `NULL` buffer, as the user may not have attached the plugin to two buffers.
`Imbuf *out`
: The image buffer for the pluginâ← s output.
`Imbuf *use`
: A pointer to the third image buffer the plugin is linked to. Plugins using this buffer should check for a `NULL` buffer, as the user may not have attached the plugin to three buffers.

`ImBuf` image structure
The `ImBuf` structure always contains 32 bits RGBA pixel data. `ImBuf` structs are always equal in size, indicated by the passed `x` and `y` value.

User Interaction

There is no way for Blender to know how many inputs a plugin expects, so it is possible for a user to attach only one input to a plugin that expects two. For this reason it is important to always check the buffers your plugin uses to make sure they are all valid. Sequence plugins should also include a text label describing the number of inputs required in the buttons interface.

Generic Sequence Plugin

```
#include "plugin.h"

char name[24]= "";

/* structure for buttons,
 *      {butcode, name,       default, min, max, tooltip}
 */
VarStruct varstr[]= {
{LABEL,   "In: X strips", 0.0, 0.0, 0.0, ""},
};

/* The cast struct is for input in the main doit function
 * Varstr and Cast must have the same variables in the same order
 */
typedef struct Cast {
int dummy; /* because of the 'label' button */
} Cast;

/* cfra: the current frame */
float cfra;

void plugin_seq_doit(Cast *, float, float, int, int, ImBuf *, ImBuf *, ImBuf *, ImBuf *);

int plugin_seq_getversion(void) {
return B_PLUGIN_VERSION;
}

void plugin_but_changed(int but) {
}

void plugin_init() {
}

void plugin_getinfo(PluginInfo *info) {
    info->name= name;
    info->nvarstr= sizeof(varstr)/sizeof(VarStruct);
    info->cfra= &cfra;
    info->varstr= varstr;
    info->init= plugin_init;
    info->seqdoit= (SeqDoit) plugin_seq_doit;
    info->callback= plugin_but_changed;
}

void plugin_seq_doit(Cast *cast, float facf0, float facf1, int xo, int yo,
                     ImBuf *ibuf1, ImBuf *ibuf2, ImBuf *outbuf, ImBuf *use) {
char *in1= (char *)ibuf1->rect;
char *out=(char *)outbuf->rect;
}
```

Our Modifications

The first step is to come up with a game plan. What is this plugin going to do, how are the users going to interact with it. For this example we will create a simple filter that will have a slider for intensity from 0-255. If any of the R,G, or B components of a pixel in the source image are less then our chosen intensity, it will return black and alpha, otherwise it will return whatever is in the image.

Now weâ← ll copy our generic plugin to `simpfilt.c` and will fill in the gaps. Its always a good idea to add some comments. First off tell users what the plugin does, where they can get a copy, who they should contact for bugs/improvements, and any licensing restrictions on the code.

When using comments make sure you use /* */ comments (C style). The plugins are in C and some C compilers do not accept // comments (C++ style).

```
/*
Description: This plugin is a sample sequence plugin that filters out
lower intensity pixels. It works on one strip as input.
Author: Kent Mein (mein@cs.umn.edu)
Website: http://www.cs.umn.edu/~mein/blender/plugins
Licensing: Public Domain
Last Modified: Sun Sep 7 23:41:35 CDT 2003
*/
```

Next we need to fill in the `name`, you should really keep this the same as your `*.c` file. Preferably descriptive, less than 23 chars, no spaces, and all lowercase.

```
char name[24]= "simpfilt.c";
```

The Cast and varstr need to be in sync. We want one slider so weâ← ll do the following:

```
VarStruct varstr[]= {
{LABEL,   "In: 1 strips", 0.0, 0.0,   0.0, ""},
{NUM|INT, "Intensity",   10.0, 0.0, 255.0, "Our threshold value"},
};

typedef struct Cast {
int dummy; /* because of the 'label' button */
int intensity;
} Cast;
```

Now we need to fill in `plugin_seq_doit`. We basically want to loop through each pixel and if RGB are all less than intensity, set the output pixel to `(0,0,0,255)`, else set it to the input values for that position.

```
void plugin_seq_doit(Cast *cast, float facf0, float facf1, int xo, int yo,
                     ImBuf *ibuf1, ImBuf *ibuf2, ImBuf *outbuf, ImBuf *use) {
int nbrComp = xo * yo * 4; /*There are xo times yo pixels, with 4 components for each;*/
int i;
char *in1 = (char *)ibuf1->rect; /*ImBuf::rect is an int pointer (containing 4 chars!) */
char *out = (char *)outbuf->rect;

for(i=0; i < nbrComp; i+=4)
{
/* if R and G and B of the pixel are all greater than
                   the intensity, keep it unchanged. */
if((in1[i+0] > cast->intensity) &&
   (in1[i+1] > cast->intensity) &&
   (in1[i+2] > cast->intensity))
{
   [i+0] = in1[i+0];    out
   [i+1] = in1[i+1];    out
```

```
   [i+2] = in1[i+2];    out
   [i+3] = in1[i+3];    out
}
/* else, make the pixel black and transparent! */
else
{
   [i+0] = out[i+1] = out[i+2] = 0;
   [i+3] = 255;         out
}
}
}
```

So we wind up with â← simpfilt.câ← .

Compiling

â← bmakeâ← is a simple utility (shell script) to aid in the compilation and development of plugins, and can be found in the `plugins/` sub-directory of the Blender installation directory. It is invoked by: `bmake (`*`plugin_name`*`.c)` and will attempt to link the proper libraries and compile the specified C file properly for your system.

If you are trying to develop plugins on a windows machine, bmake may not work for you. In that case you should look into using â← lccâ← . You can use the following to compile a plugin with lcc:

Assuming you have your plugins in `c:\blender\plugins`, here is an example of how you would compile the sequence plugin `sweep.c`. Open a dos prompt and do the following (youâ← ll want to make sure the `lcc\bin` directory is in your path):

```
cd c:\blender\plugins\sequence\sweep
lcc -Ic:\blender\plugins\include sweep.c
lcclnk -DLL sweep.obj c:\blender\plugins\include\seq.def
implib sweep.dll
```

The latest version of lcc does not have â← implibâ← . Instead, do the following:

```
cd c:\blender\plugins\sequence\sweep
lcc -Ic:\blender\plugins\include sweep.c
lcclnk -dll -nounderscores sweep.obj c:\blender\plugins\include\seq.def
```

Note that `seq.def` is not distributed with the Blender source but is available from the Blender plugins repository. Alternatively, create a `seq.def` file, and copy and paste it the following:

```
EXPORTS
LibMain@12
plugin_but_changed
plugin_getinfo
plugin_init
plugin_seq_doit
plugin_seq_getversion
```

Blender Game Engine Overview

Blender has its own built in Game Engine that allows you to create interactive 3D applications. The Blender Game Engine (BGE) is a powerful high-level programming tool. Its main focus is Game Development, but can be used to create any interactive 3d software for other things, such as interactive 3d architectural tours or educational physics research.

Using the Game Engine

The core of the BGEâ← s structure are Logic Bricks. The goal of Logic Bricks is to offer an easy to use visual interface for designing interactive applications without any programming language knowledge. There are three types of Logic Bricks, Sensors, Controllers and Actuators. Each one is better detailed here:

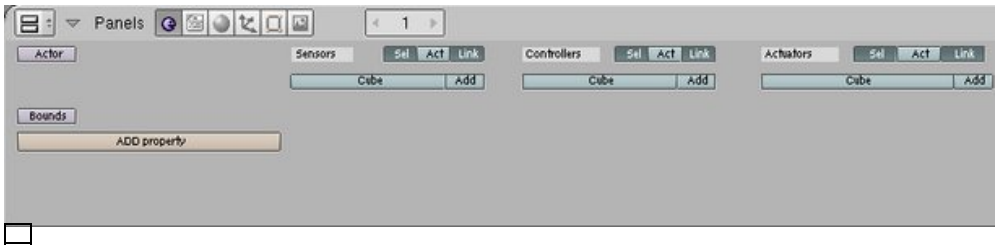
- Sensors
- Controllers
- Actuators

If you prefer to write games using Python, the game engine also has its own [Python API], separate from the rest of Blender, which you can use to write scripts to control your game. This is done by creating a Python Controller and linking it to a python script.

Game Logic

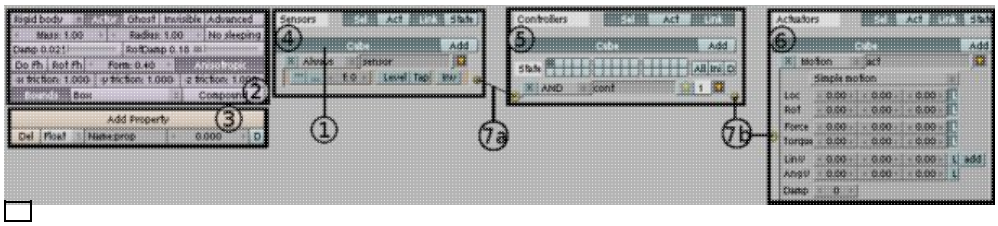
Game Logic is what causes anything to happen in the game. It is designed to provide a powerful tool to set up the logic through a graphical interface. The blocks (or â← bricksâ←) represents preprogrammed functions which can be tweaked and combined to create the game/application. The system is broken up in three parts: sensors, controllers and actuators. Sensors sense when things happens, such as a collision, a key press, mouse movement. Sensors are linked to controllers, which compare them and activate actuators.

When loading the Logic panel (context) for the first time, either by pressing F4 or selecting the logic panel button (), you will see something similar to the image below.



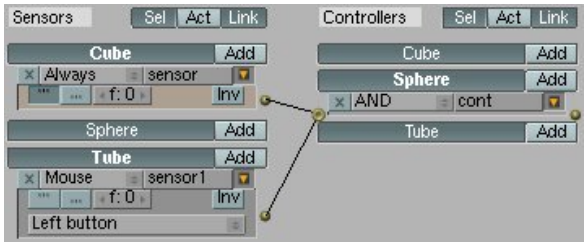
The Logic context.

To give you a better understanding of the Logic panel, we have expanded some of the menus and numbered them in the image below. We will look at each section individually as numbered in the image below.



The different parts of the Logic context.

1 Associated Object(s)



Logic for several objects.

Logic in the Blender game engine is attached to Objects. Objects are listed by their name and appear in the logic window when they are selected. You can select a single object as in the picture above, or you can select multiple objects as in the picture to the right.

An objectâ← s logic is only visible when the object(s) associated with that logic are selected.

2 Physics

These settings control how an object behaves in the game. The object type menu defines if an object is a:

- Soft body - Physical, deforms on collision.
- Rigid body - Physical.
- Dynamic - Physical, stays upright.
- Static - Physical, but unaffected by gravity.
- Sensor - For advanced sensoring options (release notes).
- No collision - Nonphysical.
- Occluder - Stops the objects behind them from being sent to the GPU, to save performance (release notes).

The other settings depend of the object type, only the relevant settings will be visible. For some types there is also an Advanced menu with more options. Some settings are per material and can be accessed in by pressing the DYN button in the Material panneau, Material sub-context.

For more physics and culling options see Mist/Stars/Physics tab, in the World sub-context.

A more in-depth look at this is available on the page Object type.

3 Properties

Properties are like variables in other programming languages. They are used to save and access data associated with an object. The following types are available:

- Timer - Starts at a defined number and counts up.
- String - Holds text.
- Float - Holds decimal numbers between -10000.000 and 10000.000.
- Integer - Holds integer numbers between -10000 and 10000.
- Bool - Is either True or False.

For a more in-depth look, see Properties.

4 Sensors

Sensors begin all logic actions. Sensors â← senseâ← things, such as a nearby object, a key pressed on the keyboard, timed events, etc. When a sensor is triggered, a pulse is sent to all linked controllers.

For a more in-depth look, see Sensors.

5 Controllers

Controllers handle the logic, evaluating pulses from sensors and sending pulses to actuators in response. The different kinds of controllers are:

- AND - All connected sensors must be positive to send a positive pulse.
- OR - One or more connected sensor has to be positive.
- XOR - Exclusive Or: one, and only one, connected sensor has to be positive.
- NAND - Not And, inverted And controller.

- NOR - Not Or, inverted Or controller.
- XNOR - Exclusive Nor controller.
- Expression - Write your own expression.
- Python - Control the sensor with a python script or module.

For a more in-depth look, see Controllers.

6 Actuators

Actuators affect objects or the game in some way. Actuators change motion, sound, properties, objects, etc. These changes can be in other objects, physics, properties or they can cause trigger events for other logic bricks.

For a more in-depth look, see Actuators.

7 Links

Links (7a) are the direction of logical flow between objects. Link lines are drawn by LMB dragging from one link node (7b) to another. Links can only be drawn from Sensors to Controllers or Controllers to Actuators. Sending nodes (the yellow circles found on the right-hand side of Sensors and Controllers) can send to multiple reception nodes (found on the left-hand side of Controllers and Actuators). Reception nodes can like wise receive multiple links.

You cannot directly link Sensors to Actuators. Actuators cannot be linked back to Sensors. If you desire to cause a Sensor to activate after an Actuator has completed, use the actuator sensor.

Object types

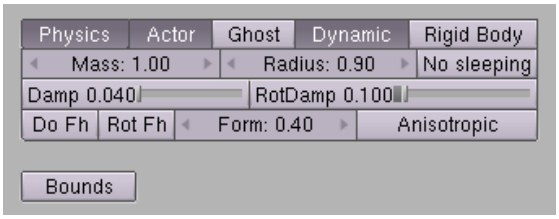


Objects type menu
You can have several object types. The Object type menu is where you define all properties associated with active objects. Active objects are evaluated by the physics engine and treated as more than simple world objects.

You can choose these object types:

- Occluder
- No collision
- Sensor
- Static
- Dynamic
- Rigid body
- Soft body

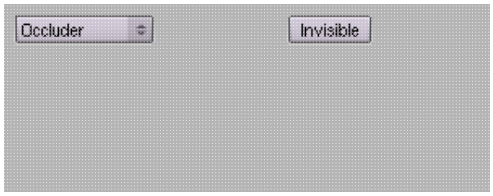
Object physics in case of no active World



Objects physics with no active world
If you don't have an active World in the Shading Context (Buttons window, F5), instead of the Object type panel you will just see a button named "Physics": enabling it, you will be able to click on more buttons to obtain a panel as in Objects physics with no active world. In this case, the available options are documented here.

In general, you need to create a world because it also sets up a physics environment.

Occluder object type



☐
Occluder
The object prevents the objects behind itself from being sent to the GPU, to save performance. (see release notes about Occlusion Culling)

Invisible

If the object is set to "Invisible", it isn't rendered in the Game Engine.

Technical specifications

(from Blender 2.49 Release Notes)

View frustrum culling

The Dynamic Bounding Volume Tree (DBVT) broadphase facility of the Bullet Physics Library is used to build a tree of graphical objects in the scene. The elements of the tree are Aabb boxes (Aligned Axis Bounding Boxes) enclosing the objects. This provides good precision in closed and opened scenes. This new culling system is enabled by default but just in case, it can be disabled with a button in the World settings.

With these improvements, the culling and Scenegraph part of the BGE is as much as 5 times faster than in 2.48. However, the speed up is noticeable only when you have many objects (>100).

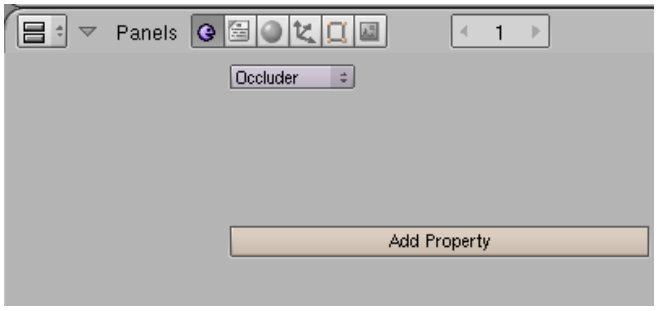
Occlusion culling

Occlusion culling is the ability of certain objects (the occluders) to hide other objects (i.e. not send them to the GPU). Just like the view frustum culling, the implementation is based on Bullet's dynamic bounding volume trees.

Occlusion culling is not enabled by default because it is a feature that requires understanding to bring the most benefits. If not used properly, it will just slow down the game. (Not too much, though, as the occlusion processing time is auto regulated.)

How it works

Occlusion culling is activated when you define at least one occluder in your scene. The Occluder option is in the physics buttons:



The "Occluder" option should only be used on a mesh object. As far as Physics is concerned, it is equivalent to "No collision". The reason why I chose to make the Occluder mode mutually exclusive with other physics mode is to emphasize on the fact that occluders should be specifically designed for that purpose and not every mesh should be an occluder. However, you can enable the Occlusion capability on physics objects using Python and Logic bricks. We'll come to that.

When an occluder object enters the view frustrum, the BGE builds a ZDepth buffer from the faces of that object. Whether the faces are one-side or two-side is important: only the front faces and two-side faces are used to build the ZDepth buffer. If multiple occluders are in the view frustrum, the BGE combines them and keeps the most foreground faces.

The resolution of the ZDepth buffer is controllable in the World settings with the "Occlu Res" button:



By default the resolution is 128 pixels for the largest dimension of the viewport while the resolution of the other dimension is set proportionally. Although 128 is a very low resolution, it is sufficient for the purpose of culling. The resolution can be increased to maximum 1024 but at great CPU expense.

The BGE traverses the DBVT tree and for each node checks if it is entirely hidden by the occluders and if so, culls the node (and all the objects it contains).

To further optimize the feature, the BGE builds and uses the ZDepth buffer only when at least one occluder is in the view frustrum. Until then, there is no performance decrease compared to regular view frustrum culling.

How to use it

There are situations where occlusion culling will not bring any benefit:

- If the occluders are small and don't hide many objects.
 In that case, occlusion culling is just dragging your CPU down.
- If the occluders are large but hides simple objects.
 In that case you're better off sending the objects to the GPU.
- If the occluders are large and hides many complex objects but in a very predictable way.
 Example: a house full of complex objects. Although occlusion culling will perform well in this case, you will get better performance by implementing a specific logic that hides/unhides the objects; for instance making the objects visible only when the camera enters the house.

Occlusion culling makes most sense when the occluders are large objects (buildings, mountains, ...) that hide many complex objects in an unpredictable way. However, don't be too concerned about performance: even if you use it inappropriately, the performance decrease will be limited due to the structure of the algorithm.

Occluders can be visible graphic objects but beware that too many faces will make the ZDepth buffer creation slow. For example, a terrain is not a good candidate for occlusion: too many faces and too many overlap. Occluder can be invisible objects placed inside more complex objects (ex: "in the walls" of a building with complex architecture). Occluders can have "holes" through which you will see objects.

Here is a possible workflow to create occluders from a complex terrain:

- duplicate the terrain and make it low poly
- remove all horizontal parts and keep only the high elevation parts
- make all the face invisible and set the Occluder mode.

To optimize further, you may split the occluder into several objects (make an occluder from each mountain, each wall, etc.) and do the same for the terrain: this way, parts of the terrain that are behind occluders will not be rendered.

You can verify how the occlusion is working by running the game in wireframe mode.

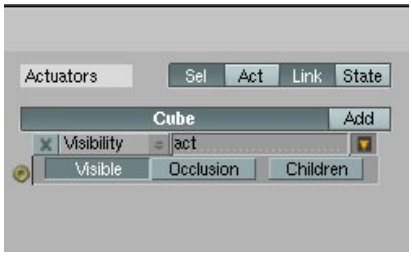
In-game control

You can turn on/off the occlusion capability of an object with Python:

```
obj.occlusion = True
obj.setOcclusion(True, False)    #param1=occlusion, param2=recurse in children
```

The object doesn't need to be defined as an occluder in the GUI: you can enable occlusion capability on any mesh object, including static and dynamic objects during the game.

You can also use the Visibility Actuator:



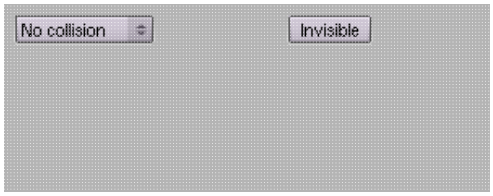
Note that the layout of the actuator has changed: The Visible/Invisible mutually exclusive buttons have been replaced by two independent buttons: push the button to activate the feature (Visible/Occluding), unpush the button to activate the opposite feature (Invisible/Non-occluding).

Note that the actuator sets the Visibility and Occlusion mode at the same time. If you need to set them separately, use Python.

Known limitation

- If you want to use occlusion culling, you must at least have one occluder at the start of the game. The occluder doesn't need to be in an active layer. After that, you can enable and disable occlusion capability on every mesh objects during the game using on the In-game control.

No collision object type

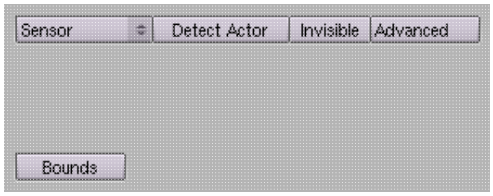


☐
No collision
This disables object collisions. Use this for objects which can never be touched by any dynamic object to save performance.

Invisible

If the object is set to "Invisible", it isn't rendered in the Game Engine.

Sensor object type



☐
Sensor
The object detects static and dynamic objects but not other collisions sensors objects.

Detect actor

Sensor should detect only the objects with Actor option enabled. Unset to detect all objects.

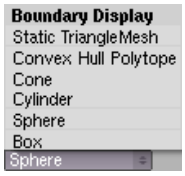
Invisible

If the object is set to "Invisible", it isn't rendered in the Game Engine.

Advanced

Collision advanced settings.

Bounds



The Blender Game Engine has a number of ways of dealing with the shapes of objects. The reason that the game engine avoids highly accurate, per face object collisions, is due to the high level of processing power that it would demand. It is often faster, during game play, to uses a simpler system such as a box or a sphere for an object thatâ ← s shape is very similar.



The bounds are calculated from the objectâ ← s centre, which is designated by the pink dot normally at the centre of your object. It is not uncommon or useful to move this centre to a more convenient location. This can be done in the Editing context (F9) and choosing one of the 3 options that focus around Centre movement (see Mesh for more info).

Convex Hull and static Triangle Mesh are designed for uses with World objects floors, walls, trees, etc.

Blender has no generic way of interacting with hollow objects. An example, which would be difficult to create, would be a cup to shake dice in.

Technical specifications

(from Blender 2.49 Release Notes)

Sensor Object Type

New sensor object to generalize Near and Radar sensor, static-static collision capability.

A new type of "Sensor" physics object is available in the GE for advanced collision management. It's called Sensor for its similarities with the physics objects that underlie the Near and Radar sensors. Like the Near and Radar object it is:

- static and ghost
- invisible by default
- always active to ensure correct collision detection
- capable of detecting both static and dynamic objects
- ignoring collision with their parent
- capable of broadphase filtering based on:
 - ♦ Actor option: the collisioning object must have the Actor flag set to be detected
 - ♦ property/material: as specified in the collision sensors attached to it.
 Broadphase filtering is important for performance reason: the collision points will be computed only for the objects that pass the broahphase filter.
- automatically removed from the simulation when no collision sensor is active on it

Unlike the Near and Radar object it can:

- take any shape, including triangle mesh
- be made visible for debugging (just use the Visible actuator)
- have multiple collision sensors using it

Other than that, the sensor objects are ordinary objects. You can move them freely or parent them. When parented to a dynamic object, they can provide advanced collision control to this object.

The type of collision capability depends on the shape:

- box, sphere, cylinder, cone, convex hull provide volume detection.
- triangle mesh provides surface detection but you can give some volume to the suface by increasing the margin in the Advanced Settings panel. The margin applies on both sides of the surface.

Performance tip:

- Sensor objects perform better than Near and Radar: they do less synchronizations because of the Scenegraph optimizations and they can have multiple collision sensors on them (with different property filtering for example).
- Always prefer simple shape (box, sphere) to complex shape whenever possible.
- Always use broadphase filtering (avoid collision sensor with empty propery/material)
- Use collision sensor only when you need them. When no collision sensor is active on the sensor object, it is removed from the simulation and consume no CPU.

Known limitations:

- When running Blender in debug mode, you will see one warning line of the console:
 "warning btCollisionDispatcher::needsCollision: static-static collision!"
 In release mode this message is not printed.
- Collision margin has no effect on sphere, cone and cylinder shape.

Other performance improvements:

- Remove unnecessary interpolation for Near and Radar objects and by extension sensor objects.
- Use direct matrix copy instead of quaternion to synchronize orientation.

Static object type



☐
Static
Physical object, but not affected by gravity.

Actor

This causes the physics engine to evaluate your object, causing your object to be seen as more than part of the world (a simple wall or floor). This object can now be seen by other active objects.

At this point your object is not affected by normal physics (e.g. gravity) but is seen as more than just a background object. This is useful for interaction between objects that need not have physics properties, such as gravity, e.g. a button on a wall.

Ghost

Ghost gives the object the ability to pass through other objects. This helps save processing time while calculating objectsâ← physics (e.g. bouncing, friction, etc.). A collision with another object is still detected and reported to your game logic.

Invisible

If the object is set to "Invisible", it isn't rendered in the Game Engine.

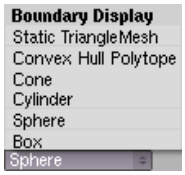
Advanced

Collision advanced settings.

Anisotropic

This controls the direction of friction in the X, Y or Z direction of the object. Anisotropic is relative to the amount of friction in a particular direction. For example, take the wheel of a skate board, the forward direction should rotate easily, yet a sideward motion should be difficult and decreased by a higher level of friction.

Bounds



The Blender Game Engine has a number of ways of dealing with the shapes of objects. The reason that the game engine avoids highly accurate, per face object collisions, is due to the high level of processing power that it would demand. It is often faster, during game play, to uses a simpler system such as a box or a sphere for an object thatâ← s shape is very similar.

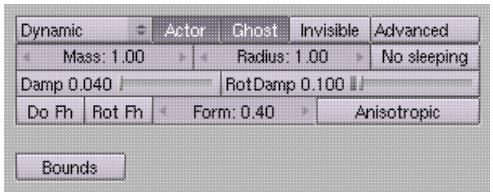


The bounds are calculated from the objectâ← s centre, which is designated by the pink dot normally at the centre of your object. It is not uncommon or useful to move this centre to a more convenient location. This can be done in the Editing context (F9) and choosing one of the 3 options that focus around Centre movement (see Mesh for more info).

Convex Hull and static Triangle Mesh are designed for uses with World objects floors, walls, trees, etc.

Blender has no generic way of interacting with hollow objects. An example, which would be difficult to create, would be a cup to shake dice in.

Dynamic object type



☐
Dynamic
Once activated, this gives your object real physics properties. Your object will now be pulled by gravity, bounce off other objects, have mass and be slowed by friction.

Actor

This causes the physics engine to evaluate your object, causing your object to be seen as more than part of the world (a simple wall or floor). This object can now be seen by other active objects.

At this point your object is not affected by normal physics (e.g. gravity) but is seen as more than just a background object. This is useful for interaction between objects that need not have physics properties, such as gravity, e.g. a button on a wall.

Ghost

Ghost gives the object the ability to pass through other objects. This helps save processing time while calculating objectsâ← physics (e.g. bouncing, friction, etc.). A collision with another object is still detected and reported to your game logic.

Invisible

If the object is set to "Invisible", it isn't rendered in the Game Engine.

Advanced

Collision advanced settings.

Mass

Mass affects the amount of force that is required to move an object. The greater the mass the greater the force that need to be applied to the object in order to move the object. Mass will not cause objects to fall faster or slower. To slow an object down use dampening, if you wish to accelerate the falling rate increase gravity (normal gravity is 9.81).

Radius

This determines the size of the bounding area for sphere bound objects and Fh/FhRot. This can be changed to a large or smaller size to give your object the look and feel you desire. The shape of the bounding area is determined by the Bounds button, by default it is a Box.

When changing the bounding size you can see the effect on the screen, in the shape of a circle. At this point in time all bounding shapes are reflected as a circle in the 3D window, even though they react differently during the game play.

No Sleeping

By default if an object that is affected by rolling physics comes to a rest, Rolling Physics is disabled. The physic engine assumes that it no longer needs to calculate the objectâ← s motion. In some case this is true and will save you processing power, in other case such as in a soccer game, the physics needs to be maintained until someone kicks the ball placing it back in motion.

Damp

Dampening affects the freedom of an object to move. In space there should be almost no dampening, while in water the dampening should be set quite high.

RotDamp

Rotational Dampening is the same as dampening but only affects the objectâ← s rotation. A high Rotational Dampening will not affect the rest of the objectâ← s momentum.

Do Fh

This causes your object to interact with Dynamic Materials. This is currently not supported by the default bullet physics and is not fully supported in the Sumo engine in the current edition of Blender. It is fully supported by version 2.24 of Blender.

Rot Fh

This option causes the Z axis to be pointed away from a Dynamic Materials. This is activated by Do Fh.

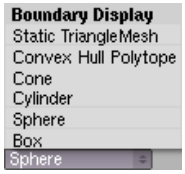
Form

The Form factor gives you a degree of control over rigid body objects. The higher the Form factor, the less likely an object is to roll, especially rolling on a flat surface.

Anisotropic

This controls the direction of friction in the X, Y or Z direction of the object. Anisotropic is relative to the amount of friction in a particular direction. For example, take the wheel of a skate board, the forward direction should rotate easily, yet a sideward motion should be difficult and decreased by a higher level of friction.

Bounds



The Blender Game Engine has a number of ways of dealing with the shapes of objects. The reason that the game engine avoids highly accurate, per face object collisions, is due to the high level of processing power that it would demand. It is often faster, during game play, to uses a simpler system such as a box or a sphere for an object thatâ ← s shape is very similar.

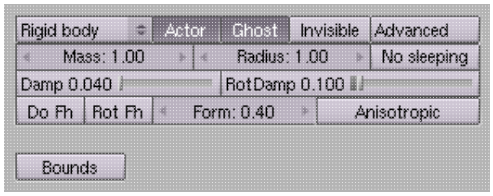


The bounds are calculated from the objectâ ← s centre, which is designated by the pink dot normally at the centre of your object. It is not uncommon or useful to move this centre to a more convenient location. This can be done in the Editing context (F9) and choosing one of the 3 options that focus around Centre movement (see Mesh for more info).

Convex Hull and static Triangle Mesh are designed for uses with World objects floors, walls, trees, etc.

Blender has no generic way of interacting with hollow objects. An example, which would be difficult to create, would be a cup to shake dice in.

Rigid body object type



☐
Rigid body
This enables rolling physics, giving objects a life-like feel. If a ball drops on a level surface naturally, it would seem strange if it just came to a halt. With rolling physics enabled your ball even dropping on a perfectly level surface and will roll as a natural object would.

Actor

This causes the physics engine to evaluate your object, causing your object to be seen as more than part of the world (a simple wall or floor). This object can now be seen by other active objects.

At this point your object is not affected by normal physics (e.g. gravity) but is seen as more than just a background object. This is useful for interaction between objects that need not have physics properties, such as gravity, e.g. a button on a wall.

Ghost

Ghost gives the object the ability to pass through other objects. This helps save processing time while calculating objectsâ← physics (e.g. bouncing, friction, etc.). A collision with another object is still detected and reported to your game logic.

Invisible

If the object is set to "Invisible", it isn't rendered in the Game Engine.

Advanced

Collision advanced settings.

Mass

Mass affects the amount of force that is required to move an object. The greater the mass the greater the force that need to be applied to the object in order to move the object. Mass will not cause objects to fall faster or slower. To slow an object down use dampening, if you wish to accelerate the falling rate increase gravity (normal gravity is 9.81).

Radius

This determines the size of the bounding area for sphere bound objects and Fh/FhRot. This can be changed to a large or smaller size to give your object the look and feel you desire. The shape of the bounding area is determined by the Bounds button, by default it is a Box.

When changing the bounding size you can see the effect on the screen, in the shape of a circle. At this point in time all bounding shapes are reflected as a circle in the 3D window, even though they react differently during the game play.

No Sleeping

By default if an object that is affected by rolling physics comes to a rest, Rolling Physics is disabled. The physic engine assumes that it no longer needs to calculate the objectâ← s motion. In some case this is true and will save you processing power, in other case such as in a soccer game, the physics needs to be maintained until someone kicks the ball placing it back in motion.

Damp

Dampening affects the freedom of an object to move. In space there should be almost no dampening, while in water the dampening should be set quite high.

RotDamp

Rotational Dampening is the same as dampening but only affects the objectâ← s rotation. A high Rotational Dampening will not affect the rest of the objectâ← s momentum.

Do Fh

This causes your object to interact with Dynamic Materials. This is currently not supported by the default bullet physics and is not fully supported in the Sumo engine in the current edition of Blender. It is fully supported by version 2.24 of Blender.

Rot Fh

This option causes the Z axis to be pointed away from a Dynamic Materials. This is activated by Do Fh.

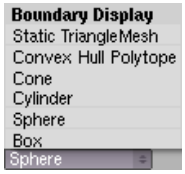
Form

The Form factor gives you a degree of control over rigid body objects. The higher the Form factor, the less likely an object is to roll, especially rolling on a flat surface.

Anisotropic

This controls the direction of friction in the X, Y or Z direction of the object. Anisotropic is relative to the amount of friction in a particular direction. For example, take the wheel of a skate board, the forward direction should rotate easily, yet a sideward motion should be difficult and decreased by a higher level of friction.

Bounds



The Blender Game Engine has a number of ways of dealing with the shapes of objects. The reason that the game engine avoids highly accurate, per face object collisions, is due to the high level of processing power that it would demand. It is often faster, during game play, to uses a simpler system such as a box or a sphere for an object thatâ← s shape is very similar.

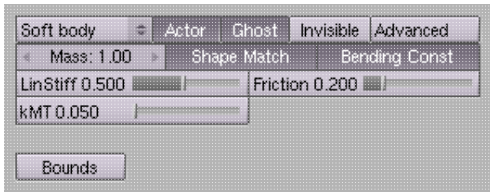


The bounds are calculated from the objectâ← s centre, which is designated by the pink dot normally at the centre of your object. It is not uncommon or useful to move this centre to a more convenient location. This can be done in the Editing context (F9) and choosing one of the 3 options that focus around Centre movement (see Mesh for more info).

Convex Hull and static Triangle Mesh are designed for uses with World objects floors, walls, trees, etc.

Blender has no generic way of interacting with hollow objects. An example, which would be difficult to create, would be a cup to shake dice in.

Soft body object type



☐
Soft body
Physical object, deforms on collision.

Actor

This causes the physics engine to evaluate your object, causing your object to be seen as more than part of the world (a simple wall or floor). This object can now be seen by other active objects.

At this point your object is not affected by normal physics (e.g. gravity) but is seen as more than just a background object. This is useful for interaction between objects that need not have physics properties, such as gravity, e.g. a button on a wall.

Ghost

Ghost gives the object the ability to pass through other objects. This helps save processing time while calculating objectsâ← physics (e.g. bouncing, friction, etc.). A collision with another object is still detected and reported to your game logic.

Invisible

If the object is set to "Invisible", it isn't rendered in the Game Engine.

Advanced

Collision advanced settings.

Mass

Mass affects the amount of force that is required to move an object. The greater the mass the greater the force that need to be applied to the object in order to move the object. Mass will not cause objects to fall faster or slower. To slow an object down use dampening, if you wish to accelerate the falling rate increase gravity (normal gravity is 9.81).

Shape match

Enables softbody shape matching goal

Bending Const

Enables bending contraints

LinStiff

Linear stiffness of the softbody links

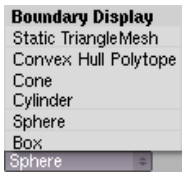
Friction

Dynamic friction

kMT

Shape matching threshold

Bounds



The Blender Game Engine has a number of ways of dealing with the shapes of objects. The reason that the game engine avoids highly accurate, per face object collisions, is due to the high level of processing power that it would demand. It is often faster, during game play, to uses a simpler system such as a box or a sphere for an object thatâ ← s shape is very similar.



The bounds are calculated from the objectâ ← s centre, which is designated by the pink dot normally at the centre of your object. It is not uncommon or useful to move this centre to a more convenient location. This can be done in the Editing context (F9) and choosing one of the 3 options that focus around Centre movement (see Mesh for more info).

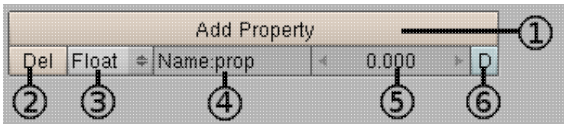
Convex Hull and static Triangle Mesh are designed for uses with World objects floors, walls, trees, etc.

Blender has no generic way of interacting with hollow objects. An example, which would be difficult to create, would be a cup to shake dice in.

Properties

Properties are the game logic equivalent to variables. They are stored on the object and can be used to represent things about them such as ammo, health, name, and so on.

Properties Panel



The properties panel.
The properties panel is made up six parts:

(1) Add Property button
: This button adds a new property to the list, default is a Float property named â← `prop`â← , followed by a number if there already is one with this name.

(2) Delete button
: Next to each property is the button to delete it.

(3) Property type menu
: It determines which type of property it is (see bellow).

(4) Name field
: Where you give your property its name, this is how you are going to access it through python or expressions. The way to do so in python is by dictionary style lookup (`GameObject["`*`propname`*`"]`). The name is case sensitive.

(5) Value field
: Sets the initial value of the property.

(6) Debug button
: Turns on debugging. Note that you must also tick Show Debug Properties in the Game menu. When done all properties with debugging activated will be presented with their object name, property name and value during gameplay. This is useful if you suspect something with your properties is causing problems.

Property types

There is five types of properties:

Timer
: Starts at the property value and count upwards as long as the object exists. It can for example be used if you want to know how long time it takes the player to complete a level.

Float
: Uses a decimal number as value, can range from -10000.000 to 10000.000. It is useful for precision values.

Int
: Uses integers (whole numbers) as value, between -10000 and 10000. Useful for counting things such as ammunition where decimals are unnecessary.

String
: Takes text as value. Can store 128 characters.

Bool
: Boolean variable, has two values true or false. This is useful for things that has only two modes, like a light switch.

Property Sensor

Property sensor has four modes:

Equal
: Triggers a positive pulse when the property value matches the value in the sensor.

Not Equal
: Triggers a positive pulse when the property value differs from the value in the sensor.

Interval
: Triggers a positive pulse when the value of the property is between the Min and Max values in the sensor.
For â← more thanâ← , enter the property name in the Max field and the lowest number for triggering in the Min field.
For â← less thanâ← , enter the property name in the Min field and the maximum value in the Max field.
Names of other properties can also be entered to compare properties.

Changed
: Triggers a positive pulse when the value of the property changes.

See Sensors for more general information on the use of sensors and property sensor for more detailed information about the Property sensor.

Property Actuator

The Property actuator has four modes:

Assign
: Gives the the value in the Value field to the property in the Prop field. Can only be a property of the same object.

Add
: Increases the value of the property, enter a negative value to subtract. For Bool, a value other 0 (also negative) is counted as

True.

Copy
: Copies a property from another object to a property of the actuator owner.

Toggle
: Switches 0 to 1 and any other number than 0 to 0. Useful for on/off switches. Does not work with String properties.

See Actuators for more general information on the use of actuators and Property Actuator for more detailed information about the Property actuator.

Sensors

Sensors are the causes of logic to do anything. This can be a trigger event such as a nearby object, a key pressed on the keyboard, timed events, etc. When a sensor is triggered, a positive pulse is sent to all controllers that are linked to it.

Sensor viewing options

Blender sensors can be found in the Logic Panel. At the top of the Sensors sub-panel, there are four buttons labeled Sel, Act, Link and State. There is also a menu that is activated by clicking Sensors, the title of the sensor sub-panel. This is for showing/hiding sensors and is very useful so unneeded sensors can be hidden and the necessary ones visible and easier to reach. Sensors are grouped by object with every selected object appearing in a list and under each object are it's sensors.

The menu labeled Sensors looks like a label but is a menu. It has four options:

- Show Objects expands all objects.
- Hide Objects collapses all objects to just a bar with their name.
- Show Sensors expands all sensors.
- Hide Sensors collapses all sensors to bars with their names.

Both these can be controlled individually on the sensor or object.

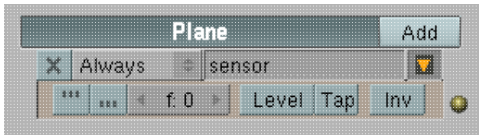
It is also possible to filter which sensors are viewed:

- Sel shows all selected objects sensors.
- Act shows only active objects sensors.
- Link shows sensors which have a link to a controller.
- State: only sensors connected to a controller of the current state are shown.

General object options



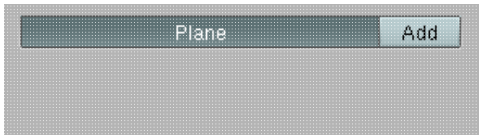
Object before adding a sensor
All objects have two buttons, one which is labeled with their name, and the other labeled Add.



Object after adding a sensor
The Add button adds a new sensor to the object.



Object with sensors collapsed
The button with the objectâ← s name is for collapsing the object list and hiding its sensors.

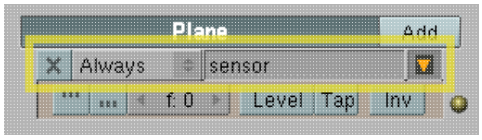


Object after deleting all sensors
An object without sensors is automatically collapsed. Adding a sensor will automatically show the object.

General sensor options

All sensors have a set of buttons and fields and menu in common. They are organized in the two rows: the **sensor header** and the **sensor pulses bar**.

Sensor header

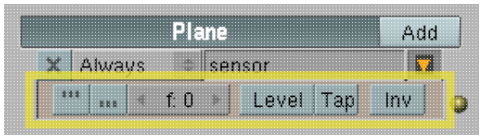


Sensor header

Cross button
: deletes the sensor.
Sensor type menu
: see sensor types below.
Sensor name
: the name of the sensor. It is used to access sensors with python; it needs to be unique among the selected objects.
Triangle button
: used for collapsing the sensor.

Sensor pulses bar

Sensors fire pulses to controllers. The pulses can be TRUE or FALSE. Different parameters on the sensor's logic block control when a sensor fires what pulse. This is extremly important as it has direct influence on the game logic and the performance. It is up to the controllers how to react on TRUE and FALSE pulses.



Sensor pulses bar

True level triggering

If this is set, the controllers receive TRUE pulses as long as the sensor's state is positive. The sensor will fire TRUE pulses with the frequency of the sensor.

False level triggering

If this is set, the controllers receive FALSE pulses as long as the sensor's state is negative. The sensor will fire FALSE pulses with the frequency of the sensor.

Note about triggers

If you don't set any triggers, the sensor fires no pulse at all as long the sensor's state does not change. When the sensor changes it's state from negative to positive the sensor fires one TRUE pulse to the controllers. When sensor changes it's state from positive to negative the sensor fires one FALSE pulse to the controllers. In between the controllers might still request the sensor's state, but if the controller does not get pulses (TRUE or FALSE) from any other sensor it will not be activated at all.

Frequency (f)

Despite the name, "Frequency", the **f** parameter sets the delay between repeated pulses, measured in logic ticks. The default value is 0 and it means no delay.

Logic ticks have a frequency of 60 Hz (60 ticks per second). For example:

- ◊ setting f=1 means the sensor pulses once every 60th of a second. A setting of f = 1, effectively means a 1 to 1 ratio between ticks, 1 tick = 1 pulse.
- ◊ setting f=30 means the sensor pulses once after 30 ticks have elapsed. This means the pulse is emitted every half a second because there are 60 ticks per second by default.
- ◊ setting f=60 means the sensor pulses every 60 ticks, which means one time per second

Raising the value of f is good for saving performance by not doing things more often than necessary.

Level

Makes the sensor working over state changes. See this page for the state system.

Tap

Sends a positive pulse only once even if the sensor remains true. Only one of Tap or Level can be activated at the same time.

When the **Tap** parameter is set the sensor will fire a FALSE pulse within the next frame, even when the sensor event is still present. When the sensor's event goes away no pulse will be fired.

If the *TRUE level triggering* is set, the TRUE/FALSE pulse pair will be repeated until the sensor's event goes away.

The *FALSE level triggering* will be ignored when the *Tap* parameter is set.

Pulses will not be inverted when the *Inv* parameter is set. But the TRUE/FALSE pulse pair will be send when the sensor's event is not present.

Inv

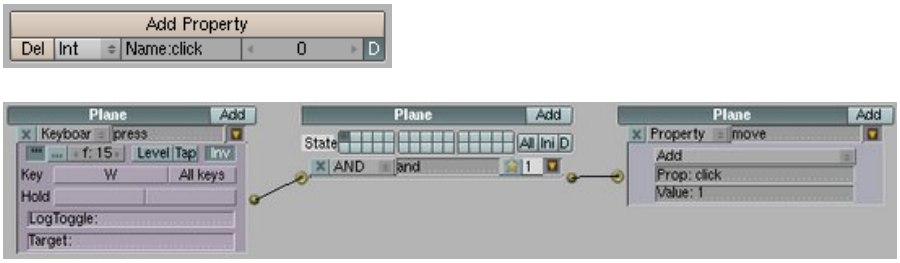
This inverts the sensor output.

If this is set, the sensor will send FALSE pulses when the sensor should send TRUE pulses and TRUE pulses if the sensor should send FALSE pulses. If the *Tap* parameter is set, the sensor acts different (refer to the previous section).

Note about Inv and triggers

Note that the toggle Inv inverts the level BEFORE the triggers, which means the triggers act on the signal coming out from Inv.

As a test, setup the logic bricks as below on an object of your choice.

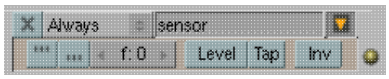


Start the game and and don't press W: in the debug info you should see the "click" property running, changing every 15 ticks (a quarter of a second). This is because the sensor is FALSE but it is inverted to TRUE with Inv, hence the True level trigger is pulsing.

GameLogic python API.

Sensor types

Always sensor

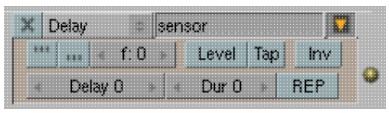


Always sensor
The Always sensor is used for things that need to be done every logic tick, or at every *x* logic tick (with non-null f), or at start-up (with Tap).

This sensor doesnâ ← t have any options apart from the general sensor ones.

GameLogic python API.

Delay sensor



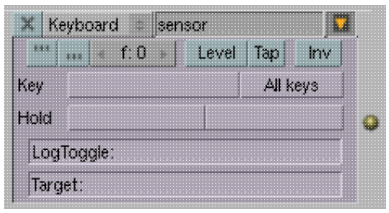
Delay sensor
The Delay sensor is designed for delaying reactions a number of logic ticks. This is useful if an other action has to be done first or to time events.

The Delay sensor has three options:

- Delay is the number of logic ticks the sensor waits before sending a positive pulse.
- Duration is the time the sensor waits before sending the negative pulse.
- Repeats makes the sensor restart after the delay and duration time is up.

GameLogic python API.

Keyboard sensor



Keyboard sensor
The Keyboard sensor is for detecting keyboard input, it can also save keyboard input to a String property.

The first row of buttons (Key one) is for single key presses. Press the button with no label and a key to assign that key to the sensor. This is the active key, which will trigger the positive pulse. Click the button and then click outside of the button to deassign the key. The All keys button will hide all other button and send positive

pulses on all key presses. This is useful for custom key maps with a Python controller.

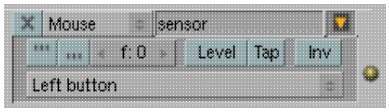
The Hold buttons work the same way as the key button but wonâ ← t set off a positive pulse unless hold down while the key button is pressed. You can have up to two of hold buttons. If a key is assigned here, a positive pulse wonâ ← t be send unless both/all keys are pressed. This is useful if you wish to have for example CtrlR or â § ShiftAltEsc to do a specific action.

LogToggle field assigns a Bool property which determines if the keystroke will be logged in the String property (true) or not (false). This needs to be here if you wish to log your keystrokes.

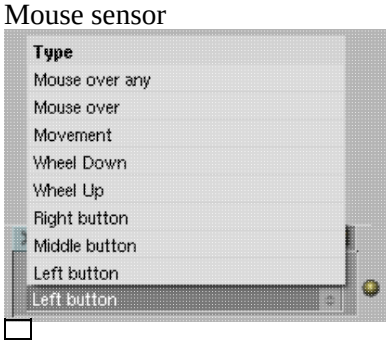
Target is the String property to which the keystrokes are saved. Together with a Property sensor this can be used for example to enter passwords.

GameLogic python API.

Mouse sensor



Mouse sensor



Mouse sensor input type

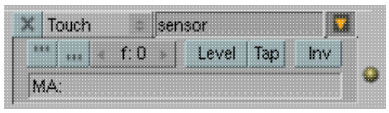
The Mouse sensor is for detecting mouse input. The controller consist only of a list of types of mouse events. These are:

- Left button.
- Middle button.
- Right button
- Wheel Up, the scroll wheel of the mouse.
- Wheel Down, the scroll wheel of the mouse.
- Movement, any movement with the mouse.
- Mouse over, detects if the mouse is over the object.
- Mouse over any, detects if the mouse is over any object.

There is no logic brick for specific mouse movement and reactions (such as first person camera), these has to be coded in python.

GameLogic python API.

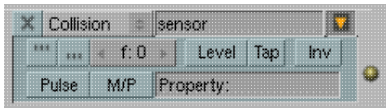
Touch sensor

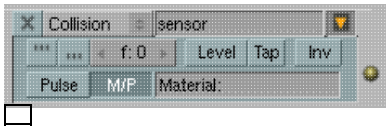


Touch sensor
The Touch sensor sends a positive pulse when the object is in contact with another object. The MA field is for filtering materials. Only contact with the material in this field will generate a positive pulse. Leave blank for touch with any object. The positive pulse is sent on collision and the negative pulse is sent once the objects are no longer in contact. For a continuous pulse while they are in contact use â← True Pulse triggeringâ← .

GameLogic python API.

Collision sensor



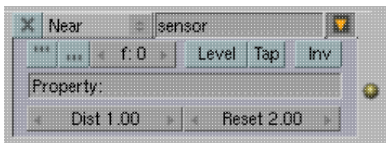
Collision sensor


Collision sensor
A Collision sensor works like a Touch sensor but can also filter by property. Only objects with the property with that name will generate a positive pulse upon collision. Leave blank for collision with any object.

The Pulse button makes it sensible to other collisions even if it is still in touch with the object that triggered the last positive pulse.

The M/P button toggles between material and property filtering.

Near sensor

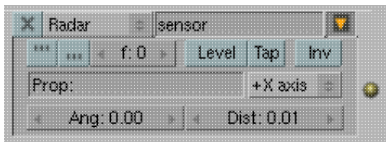


Near sensor
A Near sensor detects objects that are within a specific distance of themselves. It can filter objects with properties, like the Collision sensor.

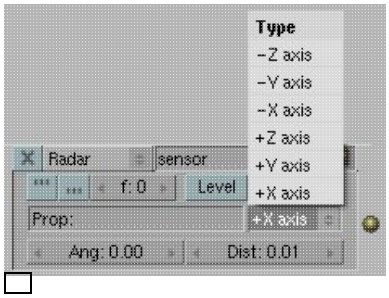
Distance is the number of blender units it will detect objects within. Reset is the distance the object needs to be to reset the sensor (send a negative pulse).

GameLogic python API.

Radar sensor



Radar sensor



Radar sensor axis menu

A Radar sensor works much like a Near sensor, but only within an angle from an axis, forming an invisible cone with the top in the objectsâ← centre and base at a distance on an axis. It has a property filter field (Prop). Next to it is the axis menu, which determines the direction of the radar cone. The Â± signs is whether it is on the axis direction (+), or the opposite (-).

Angle determines the width of the cone, and Distance its length.

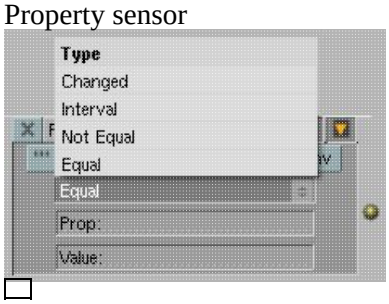
This sensor is useful for giving bots sight only in front of them, for example. Note that it does see through other objects.

GameLogic python API.

Property sensor



Property sensor



Property sensor operation type

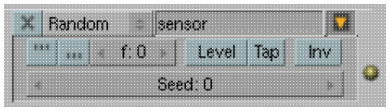
The Property sensor detect changes in the objects properties. This sensor has four modes:

- Equal triggers a positive pulse when the property value matches the value in the sensor. The Prop field is for the name of the property, and Value is for the value it has to match to send a positive pulse.

- Not Equal triggers a positive pulse when the property value differs from the value in the sensor. The same fields as Equal, but Value represents the value the property has to differ from in order to set off a pulse.
- Interval triggers a positive pulse when the value of the property is between the Min and Max values of the sensor. For â← more thanâ← , enter the property name in the Max field and the lowest number for triggering in the Min field. For â← less thanâ← , enter the property name in the Min field and the maximum value in the Max field. Names of other properties can also be entered to compare properties.
- Changed sends a positive pulse as soon as the property value changes.

GameLogic python API.

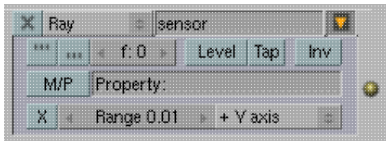
Random sensor



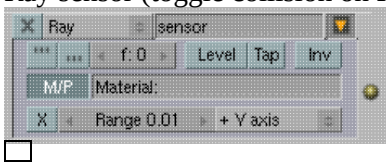
Random sensor
The Random sensor generates random pulses. It has a Seed field to enter the initial seed. 0 is not random, for testing and debugging purposes.

GameLogic python API.

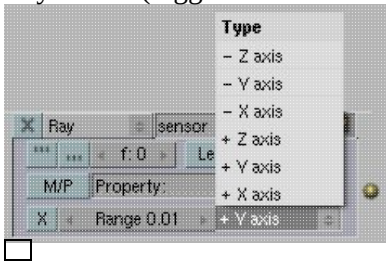
Ray sensor



Ray sensor (toggle collision on Property)



Ray sensor (toggle collision on Material)

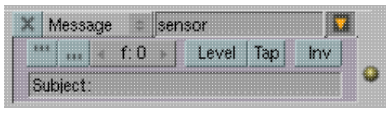


Ray sensor axis menu
The Ray sensor shoots a ray in the direction of an axis and sends a positive pulse once it hits something. It can be filtered to only detect objects with a given material or property.

It shares a lot of buttons and fields with Radar sensor. The X button make it x-ray, it sees through objects that doesnâ← t have the property or material specified in the filter field.

GameLogic python API.

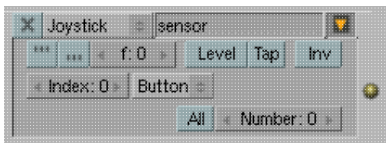
Message sensor



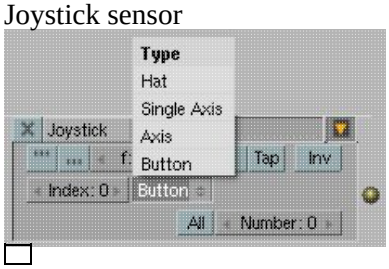
Message sensor
Messages can be used to send either text or property values. The Message sensor sends a positive pulse once a message is sent anywhere in the engine. It can filter to only send a pulse upon a message with a specific subject.

GameLogic python API.

Joystick sensor



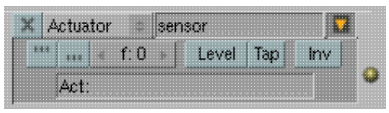
Joystick sensor



Joystick sensor events menu
The Joystick sensor detect joystick events.

GameLogic python API.

Actuator sensor



Actuator sensor
The Actuator sensor is set off when an actuator with the name specified in the Act field is activated.

GameLogic python API.

Controllers

The controllers are the bricks that collect data sent by the sensors. There are eight ways to process the input:

- AND
- OR
- XOR
- NAND
- NOR
- XNOR
- Expression
- Python

When a sensor is activated it sends out a positive pulse and when it is deactivated it sends out a negative pulse. The controllersâ← job is to check and combine these pulses to trigger the proper response.

Controller Comparison table

This table is supposed to give a quick overview of the controller types. The first column, input, represents the number of positive pulses sent from the connected sensors. The following columns represent each controllerâ← s response to those pulses. True meaning the conditions of the controller is fulfilled and the actuators it is connected to will be activated, false meaning the controllerâ← s conditions are not met and nothing will happen. Please consult the sections further down for a more detailed description of each controller.

Note
It is assumed that more than one sensor is connected to the controller. For only one sensor, consult the â← Allâ← line.

Positive sensors	**Controllers**					
	AND	OR	XOR	NAND	NOR	XNOR
None	False	False	False	True	True	True
One	False	True	True	True	False	False
Multiple, not all	False	True	False	True	False	True
All	True	True	False	False	False	True

AND Controller

The AND controller is the default type when you create a new controller. This is because it is the most common controller and it works well for just pass a sensor directly to an actuator (which cannot be done without a controller).

An AND controller activates the actuators it is connected to only if all the sensors connected to it are positive. If only one sensor is connected to it, the actuators will be activated as soon as the sensor is triggered. If more sensors are connected they all have to be activated at the same time. This is where the name comes from. If there are two sensors, â← `sensor1`â← and â← `sensor2`â← , connected to the controller both `sensor1` *AND* `sensor2` have to be active at the same time.

Example

You want to make a menu button. When the player clicks the button he is transferred to the next scene.

On the menu button you set up two Mouse sensors, a Mouse over sensor and a Left button sensor. Both are connected to the same AND controller, which is connected to a Set Scene actuator. This means the player has to both hover the mouse over the menu button and click to get to the next scene.

OR Controller

An OR controller activates the actuators it is connected to if at least one of its sensors is activated. If just one sensor is connected to the controller it will pass on the positive pulse when that sensor is activated, if more sensors are connected, only one of them need to be activated. OR is not exclusive, which means it works like an AND controller too: if all sensors are activated the OR controller will still pass the pulse. XOR controller is an exclusive or controller.

Example

You want both the Esc and the Q keys to quit your game. You set up two Keyboard sensors, one with Esc key and one with Q key. Both are then connected to an OR controller which is connected to an Quit this game actuator. If one of the two sensors is activated the game will then quit.

XOR Controller

XOR controller is an Exclusive OR controller. Exclusive means that one, *and only one*, of the sensors has to be positive. If there is only one sensor connected to the controller it will activate its actuators when that sensor sends a positive pulse. If more sensors are connected the controller will activate its actuators only when one, no matter which, of the sensors is positive. If both sensors are positive, it will not activate.

Example

---UNDER CONSTRUCTION---

NAND Controller

NAND is short for Not And. In logic, *not* negates the outcome, Not And is therefore the opposite of an AND controller. An AND controller activates actuators only when all sensors are triggered, a NAND controller does the opposite. Whenever **not** all the sensors are true, actuators are activated. If only one sensor is connected to the controller, actuators are activated when the sensor is not triggered (behaving like a simple NOT logic gate).

Example

---UNDER CONSTRUCTION---

NOR Controller

NOR is a Not OR controller, it is the opposite of an OR Controller. A NOR Controller will only activate actuators if none of the sensors are triggered.

Example

---UNDER CONSTRUCTION---

XNOR Controller

XNOR is an Exclusive Not OR controller. While sounding confusing it just has the opposite output of a XOR controller. XNOR activates actuators when more or less than one sensor is triggered.

Example

---UNDER CONSTRUCTION---

Expression Controller

A confusing controller thatâ← s very functional in special cases.




Expression controller. Property called â← `prop`â← .
In this case the Expression controller checks that the sensor's signal, called â← `sensor`â← , is true and that the value of a property called â← `prop`â← . is equal to 1.

---UNDER CONSTRUCTION---

For more helpful information on using an expression controller look here

Python Controller

The Python controller is a controller that checks the input using a user-programmed script, a.k.a. a Python script or any other file containing Python code. Python controllers have two modes: Script or Module. Both can be written in the text editor and stored inside the .blend file, or they can be external script files.

More information on Python inside BGE can be found here.

The API for Python in the BGE can be found here.

Script Mode

In Script mode, the controller is linked to a script, the whole script will be executed before the logic moves on. Everything in the script will happen at the same frame, if the same value or attribute are changed more than once, only the latest will be visible in the game because of this. If for example the position of an object is changed first to (100.0,100.0,100.0) but later in the same script changed to (0.0,0.0,0.0), then only the (0.0,0.0,0.0) will be visible to the player because the move to (100.0,100.0,100.0) happened on at the same frame.

Module Mode

Since Blender 2.49 the Python module controller was added. Simply choose Module instead of Script on your Python controller drop-down menu. Then you define a function on that module and call that function from the controller. Instead of writing â← `myScript.py`â← on the script edit box, youâ← ll write â← `myModule.myFunc`â← on the module edit box. This function will run every time the controller is called. But other functions and variables outside of `myFunc`â← s scope will only run once. This is good for optimizing your code when you want to initiate variables only once then use them later.

The Python module controller supports any number of attributes, this means packages are supported automatically. As well as â← `myModule.myFunc`â← you can do â← `myPackage.myModule.myFunc`â← , nested packages work too, as does method calls on class instances like: â← `myPackage.myModule.myInstance.myMethod`â← . The python controller is passed to the python function as an argument for functions that take one arg.

This allows live editing of scripts, learn more about the Python module controller at:

*Thread http://blenderartists.org/forum/showthread.php?t=156672
*Video http://download.blender.org/apricot/live_bge_edit.ogv

Actuators

Actuators perform actions, such as move, create objects, play a sound. The actuators initiate their functions when they get a positive pulse from one (or more) of their controllers.

There are 18 actuators as of 2.49a:

Motion	Sets object into motion and/or rotation, there are different options from â← teleportingâ← to physically push rotate objects.
Shape Action	Handles animations stored in shape keys and animated with shape actions.
Action	Handles armature actions, this is only visible if an armature is selected.
Constraint	Constraints are used to limit objectâ← s locations, distance or rotation. These are useful for controlling the physics of the object in game.
Ipo	Controls Ipo animations, these can move, rotate, scale, change colour of objects and more.
Camera	Has options to follow objects smoothly, primarily for camera objects but any object can use this.
Sound	Used to play sounds in the game.
Property	Manipulates the objectâ← s properties, like assigning, adding or copying.
Edit Object	Edits the objectâ← s mesh, adds objects or destroys them, it can also change the mesh of an object (and soon also recreate the collision mesh).
Scene	Manage the scenes in your .blend file, these can be used as levels or for UI and background.
Random	Creates random values which can be stored in properties.
Message	Sends messages, which can be received by other objects to activate them.
CD	Plays CD music (might not make it to 2.5).
Game	Handles the entire game and can do things as restart, quit, load and save.
Visibility	Changes visibility of the object.
2D Filter	Filters for special effects like sepia colours or blur.
Parent	Can set a parent to the object, or unparent it.
State	Changes states of the object.

Motion Actuator

Motion actuator sets object into motion and/or rotation, there are different options from â← teleportingâ← to physically push-rotate objects.

There are two ways of controlling motions, simple motion and servo control. Simple motion applies different kinds of motions directly while servo motion set a target speed and how quickly it reaches that speed.

The controls differ depending on what kind of physics settings the object has. Objects that are unaffected by gravity (Static, No collision, Occluder, Sensor) can only use two kinds of simple motion. The other kinds of physics objects (Dynamic, Rigid body, Soft body) can use servo control and all simple motion options.

Simple Motion

Simple motion applies motions on axes. The simple motion block is organized in rows and columns. Rows for each type of motion, columns for axes and local on/off button. The three grey boxes on each row is the motion on the X,Y,Z axis (in that order), if you want a diagonal motion, add motions to more axes at the same time. The L button switches whether the motion will be added to the local or global axis. Global axis is always the same, as shown in the 3D window with red, green, blue arrows and lines. Global axes make the object always move in the same direction no matter which way it is rotated. Local axis are the objects own axes. This is the most common option and is default.

Rotation is done around the axis entered, clockwise if you are looking along the axis direction.

The types of simple motion are:

Location (Loc)
: â← Teleportsâ← the objects. The object jumps the number of blender units entered, this can cause objects to go through other objects since it never passes the any of the coordinates between the start and end. Using the global settings will move the object on the global axis but not in global coordinates, if you wish to teleport an object to a coordinate it is recommended to use Python or Ipo.

Rotation (Rot)
: Changes the objectâ← s orientation, in the same manner as Locationâ← s changes moves it. As soon as activated the object will be turned. One revolution is represented by the value 7.2 (i.e. 0.02 for one degree).

Force
: Force is like giving the object a shove, it adds motion in a direction. But since it is added, the object will go faster and faster as long as Force is applied.

Torque
: Works like Force, but for rotation.

Linear Velocity (LinV)
: Sets a velocity, the objects move in the entered speed along the axis. The add button makes the speed added to the existing speed instead. Linear velocity is interfering with gravity, object will fall slower, this can be solved by using servo control.

Angular Velocity (AngV)
: Works like Linear Velocity, but for rotation.

Damp
: Slows the movement.

These types can be mixed in any way.

Servo Control

Servo control is a powerful way to achieve motion in the physical world. It consists in a servo controller that adjusts the force on the object in order to achieve a given speed, hence the name servo control.

Achieving speed through force is a very natural method because this is how the real world works. It produces natural movements, avoids the collision problem unlike Loc and interacts correctly with gravity unlike LinV. Servo motion uses force; it is *only* applicable to Dynamic object. It should **not** be used on Rigid body as it does not control the rotation, only the speed.

At the heart of the servo control motion actuator there is a PID servo controller: it measures the speed error (i.e. the difference between the target speed and the actual speed) and updates the force based on the error, by applying a force that is proportional to the error (the P coefficient) and proportional to the integral of the error (the I coefficient). The higher the coefficients, the â← harderâ← the speed control (i.e. quick reaction); the lower the coefficients, the â← softerâ← the speed control (i.e. slow reaction, sliding effect).

Additionally you can Limit the force along each axis so that the acceleration (or braking) force is limited. To compare with the cruise control in a car, the P and I coefficients define how precise will be the control and the Limit coefficients define the power of the engine (and the brakes), and thus how quickly it will reach the target speed.

Other options in the actuator are:

Ref
: Choose an object which the actuator owner uses as a reference for movement, for moving platforms for example. If empty it will use world reference.

LinV
: Target linear velocity, servo motion works with forces but tries to work up to and keep this velocity.

Limit
: This option limits velocity on axes, once chosen more options under will be available.
 Max/Min
 : The borders of the limit, velocity will never go outside these values.

You can achieve a great variety of movement styles with the Servo control motion actuator, for example, by not limiting the force along the transversal axis, you get a strong reaction to sliding, which is equivalent to anisotropic friction.

The servo motion actuator applies forces; it is therefore not compatible with other actuators that sets the force too, i.e. a simple motion force actuator or another servo motion, although it is possible to combine two servo motion actuators that works on different axis.

Shape Action Actuator

Shape Action actuator is working with actions that are stored in the shapekeys of the mesh and put together in the shapekey editor.

Options

It has the same animation options as the Ipo actuator, except for Ping Pong.

- Play plays the animation from start to end, and then resets.
- Flipper plays from start to end, only as long as the actuator is activated. When it is not, it plays backward back to the start.
- Loop Stop plays from start to end, only as long as the actuator is activated. When it is not, the animation is paused and continued once the actuator is activated again. When it reaches the end frame the animation starts over again.
- Loop End plays the animation from start to end, and then restarts.
- Property: for as long as the actuator is activated the animation will be controlled by the property assigned in the field replacing Start and End buttons. The property assigns the frame number.
- The AC field defines which action to be used. Note that it is case sensitive, meaning â← `myaction`â← is not the same as â← `MYACTION`â← . The Continue button makes the action continue from the last frame it was at the last time the actuator was used.
- Sta and End define the starting and ending frame, each action has its own set of frames, as opposed to Ipos which only uses one timeline.
- Blendin makes a number of frames blended between the selected action and the former, for smoother transitions between actions.
- Priority: if more than one action are played at once, the one with the lowest priority will be played. Like a list, from first to last and priority is the number in the list.
- FrameProp stores the frame number to a property.

Action Actuator

The Action actuator is only visible when an armature is selected, because actions are stored in the armature.

Options

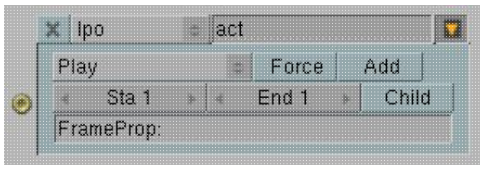
It has the same animation options as the Ipo actuator, except for Ping Pong.

- Play plays the animation from start to end, and then resets.
- Flipper plays from start to end, only as long as the actuator is activated. When it is not, it plays backward back to the start.
- Loop Stop plays from start to end, only as long as the actuator is activated. When it is not, the animation is paused and continued once the actuator is activated again. When it reaches the end frame the animation starts over again.
- Loop End plays the animation from start to end, and then restarts.
- Property: for as long as the actuator is activated the animation will be controlled by the property assigned in the field replacing Start and End buttons. The property assigns the frame number.
- The AC field defines which action to be used. Note that it is case sensitive, meaning â← `myaction`â← is not the same as â← `MYACTION`â← . The Continue button makes the action continue from the last frame it was at the last time the actuator was used.
- Sta and End define the starting and ending frame, each action has its own set of frames, as opposed to Ipos which only uses one timeline.
- Blendin makes a number of frames blended between the selected action and the former, for smoother transitions between actions.
- Priority: if more than one action are played at once, the one with the lowest priority will be played. Like a list, from first to last and priority is the number in the list.
- FrameProp stores the frame number to a property.

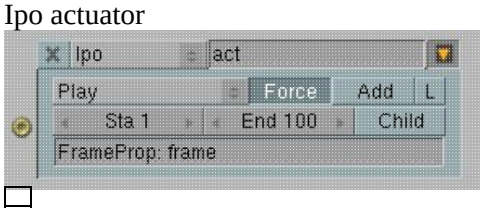
Ipo Actuator

Ipos are created with keyframes and are used to animate objects in a lot of ways.

The actuator consist of an Ipo types menu, and a set of buttons listed below:



Ipo actuator



Ipo actuator with orientation button

Force
: Applies the Ipo as force: this means that the object will not try to go through solid object. Normally it follows the Ipo curve and ignores other objects. Selecting Force makes a L button appear, for use of local/global coordinates.

Add
: Ipo curves are normally defined globally, which means the object will jump back to the position where the animation was recorded. With Add the animation is instead added to the objectâ← s position. Selecting Add makes a L button appear, for use of local/global coordinates.

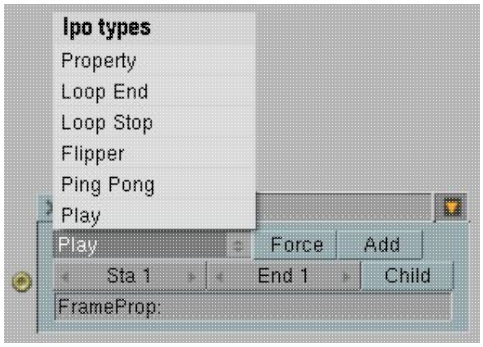
Sta (start)
: Defines which frame the animation will start at.

End
: Defines which frame is the last in the animation.

Child
: Whether childrenâ← s Ipos should be updated as well.

FrameProp
: Choose a property which will store the frame the animation is currently on through the animation.

There are six Ipo types:



Ipo actuator

Play
: Plays the animation from start to end, and then resets.

Ping Pong
: Plays the animation from start to end, next time the actuator is activated it plays from end to start.

Flipper
: Plays from start to end, only as long as the actuator is activated. When it is not, it plays backward back to the start.

Loop Stop
: Plays from start to end, only as long as the actuator is activated. When it is not, the animation is paused and continued once the actuator is activated again. When it reaches the end frame, the animation starts over again.

Loop End
: Plays the animation from start to end and then restarts.

Property
: For as long as the actuator is activated the animation will be controlled by the property assigned in the field replacing Start and End button. The property assigns the frame number.

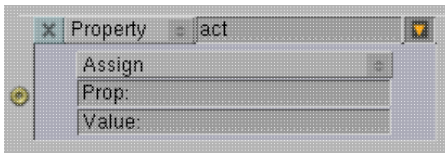
Example

You want to make a health bar which visualizes the players hp without numbers. You create a cube, move it (in Edit mode) so that the centre is in one end of the object. You stretch it out and scale it to the size you want it to be when the health is full. You decide 100 is the max hp and insert a scale keyframe at frame 100. Then you move to frame 0 and scale the object down along the right axis to zero and insert a scale keyframe. The animation will now consist of your health bar filling up between frame 0 and 100. You set up a Copy Property actuator to copy the property hp of the player to the health barâ← s property hp. Next you set up an Always sensor and AND controller and Ipo actuator in Property mode. As property you choose the hp property.

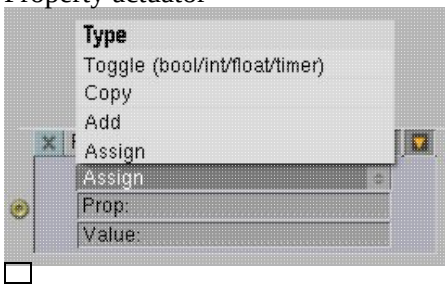
When the property hp is changed, the life bar will scale to the frame with that number. As frame 0 is no health and frame 100 is full health bar it will correspond to the property.

Property Actuator

Using the Property actuator you can change the value of a given property once the actuator itself is activated.



Property actuator



Property actuator types

Prop
: The target property that this actuator will change

Value
: The value to be used to change the property

Type

- **Assign**: the Prop target property will become equal to the set Value once the actuator is activated
- **Add**: adds Value to the value of the property Prop once the actuator is activated (enter a negative value to decrease). For Bool, a value other than 0 (also negative) is counted as True.
- **Copy**: copies a property from another object to a property of the actuator owner once the actuator is activated.
- **Toggle**: switches 0 to 1 and any other number than 0 to 0 once the actuator is activated. Useful for on/off switches.

Game Logic python API.

Example

You have a character, it has a property called â← hpâ← (hit points) to determine when he has taken enough damage to die. hp is an int with the start value of 100.

You set up two Collision sensors, one for enemy bullets, and one for picking up more health. The first one is connected (through an AND controller) to an Add Property actuator with the property hp and the value -10. Every time the player is hit by an enemy bullet he loses 10 hp. The other sensor is connected (through an AND controller) to an other Add Property actuator, this one with the value 50. So every time the player collides with a health item the hp increases by 50. Next you set up a Property sensor for an interval, greater than 100. This is connected (through an AND controller) to an Assign Property actuator which is set to 100. So if the players hp increases over 100 it is set to 100.

Edit Object Actuator

The Edit Object actuator edits the objectâ← s mesh, adds objects or destroys them, it can also change the mesh of an object (and soon also recreate the collision mesh). It has five modes:

Add Object

Adds an object at the centre of the current object. The object that is added needs to be on another, hidden, layer.

- ◊ OB: the name of the object that is going to be added.
- ◊ Time: the time (in frames) the object stays alive before it disappears. Zero makes it stay forever.
- ◊ LinV: Linear Velocity, works like in the motion actuator but on the created object instead of the object itself. Useful for shooting objects, create them with an initial speed.
- ◊ AngV Angular velocity, works like in the motion actuator but on the created object instead of the object itself.

End Object

Destroys the object.

Replace Mesh

Replaces the graphical mesh of the object. The physics is still the same. For the next version of Blender, the collision mesh will also be changeable.

- ◊ ME: the name of the new mesh.

Track To

Makes the object â← look atâ← another object, in 2D or 3D. The Y-axis is considered the front of the object.

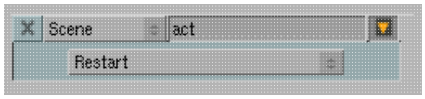
- ◊ OB: the object to look at.
- ◊ Time: the time it takes to turn to the object.
- ◊ 3D: without this it only rotate the object on the Z axis, only following the target object in two dimensions. With this it rotates on all axis following the target object in 3D space.

Dynamics

Changes physical properties of the object. It has five modes:

- ◊ Suspend Dynamics takes the object out of the physics engine, it wonâ← t respond to collisions or gravity.
- ◊ Restore Dynamics brings the object back to the physics engine.
- ◊ Disable Rigid Body makes the object not rigid body any more.
- ◊ Enable Rigid Body makes it rigid body again.
- ◊ Set Mass changes the Mass of the object.

Scene Actuator



Scene actuator
The Scene actuator manages the scenes in your .blend file, these can be used as levels or for UI and background.

The actuator has eight modes:



Scene actuator options

Restart
: Restarts the current scene, everything in the scene is reset

Set Scene
: Changes scene to selected one

Set Camera
: Changes which camera is used

Add OverlayScene
: This adds an other scene, and draws it on top of the current scene. It is good for interface, keeping the health bar, ammo meter, speed meter in an overlay scene makes them always visible

Add BackgroundScene
: This is the opposite of an overlay scene, it is drawn behind the current scene

Remove Scene
: Removes a scene!

Suspend Scene
: Pauses a scene

Resume Scene
: Resumes a paused scene.

Actuator: 2D Filters

2D Filters are image filtering actuators, that apply on final render of objects. Thanks Social et.al. for FPS_Template.

Motion Blur

Motion Blur is a 2D Filter that needs previous rendering information to produce motion effect on objects. Bellow you can see Motion Blur filter in Blender window, along with its logic bricks:



2D Filters: Motion Blur.



2D Filters: Game Logic.
To enable this filter:

1. Add appropriate Sensor(s) and Controller(s).
2. Add a 2D Filter Actuator.
3. Select Motion Blur in the drop-down list.
4. Set Motion Blur Value (Factor).

And for disabling this filter:

1. Add appropriate Sensor(s) and Controller(s).
2. Add a 2D Filter Actuator.
3. Select Motion Blur.
4. Click on D button to go to disabled mode.

You can enable Motion Blur filter using a Python controller:

```
import Rasterizer
Rasterizer.enableMotionBlur(0.85)
```

And disable it:

```
import Rasterizer
Rasterizer.disableMotionBlur()
```

Note
Your graphic hardware and OpenGL driver must support accumulation buffer (`glAccum` function).

Built-In 2D Filters

All 2D filters you can see in 2D Filter actuator have the same architecture, all built-in filters use fragment shader to produce final render view, so your hardware must support shaders.



2D Filters: Motion Blur.



2D Filters: Sepia.



2D Filters: Sobel.

Blur, Sharpen, Dilation, Erosion, Laplacian, Sobel, Prewitt, Gray Scale, Sepia and Invert are built-in filters. These filters can be set to be available in some passes.

To use a filter you should:

1. Create appropriate sensor(s) and controller(s).
2. Create a 2D Filter actuator.
3. Select your filter, for example Blur.
4. Set the pass number that the filter will be applied.

To remove a filter on a specific pass:

1. Create appropriate sensor(s) and controller(s).
2. Create a 2D Filter actuator.
3. Select Remove Filter.
4. Set the pass number you want to remove the filter from it.

To disable a filter on a specific pass:

1. Create appropriate sensor(s) and controller(s).
2. Create a 2D Filter actuator.
3. Select Disable Filter.
4. Set the pass number you want to disable the filter on it.

To enable a filter on a specific pass:

1. Create appropriate sensor(s) and controller(s)
2. Create a 2D Filter actuator.
3. Select Enable Filter.
4. Set the pass number you want to enable the filter on it.

Custom Filters



2D Filters: Custom Filter.
Custom filters give you the ability to define your own 2D filter using GLSL. Its usage is the same as built-in filters, but you must select Custom Filter in 2D Filter actuator, then write shader program into the Text Editor, and then place shader script name on actuator.

Blue Sepia Example:

```
uniform sampler2D bgl_RenderedTexture;
void main(void)
{
  vec4 texcolor = texture2D(bgl_RenderedTexture, gl_TexCoord[0].st);
  float gray = dot(texcolor.rgb, vec3(0.299, 0.587, 0.114));
  gl_FragColor = vec4(gray * vec3(0.8, 1.0, 1.2), texcolor.a);
}
```

Camera in BGE

Within the game engine, the camera may be used with its default characteristics (how it moves), or it may track an object using parent to the object, parent to a vertex of an object, or using a Camera actuator.

Note
Any object may become a camera (see below), and when using camera tracking to an object, any object may be used (such as an empty).

To use the camera you must start the GE from the cameraâ← s viewport (0 NumPad), and to avoid camera distortion, always zoom the view in until the camera object fills the entire viewport.

Default camera

It is a view of the world with no camera movement.

Parent Camera to Object

The camera will follow the object. First select the camera and then select the object. Next CtrlP â Make Parent.

Note that if your object has any rotations then the camera will also have those rotations. To avoid this use â← Parent to Vertexâ← (see below).

Parent to Vertex

The easiest way to accomplish this is to select your object and â Tab to Edit mode. Now select the vertex and â Tab back to Object mode.

Next, without any objects selected, select the camera and, holding the â § Shift key, select the object. â Tab into Edit mode, and CtrlP and choose Make vertex parent.

Now the camera will follow the object and it will maintain its rotation, while the object rotates.

Object as a Camera

Any object may also become a camera with whatever properties are set for the object.

To make an object the camera, in Object mode select the object and press Ctrl0 NumPad on the numpad.

To reverse it, just select the camera and Ctrl0 NumPad again.

See also

- Fisheye Dome camera.

WARNING
This is the copy of the Development page of the Fisheye Dome camera released in BLender 2.49. A proper manual page will be written later

Dome Mode in Blender Game Engine

This feature allows artists to visualize their interactive projects within an immersive dome environment. In order to make it an extensible tool, we are supporting Fulldome, Truncated domes (front and rear), Planetariums and domes with spherical mirrors.

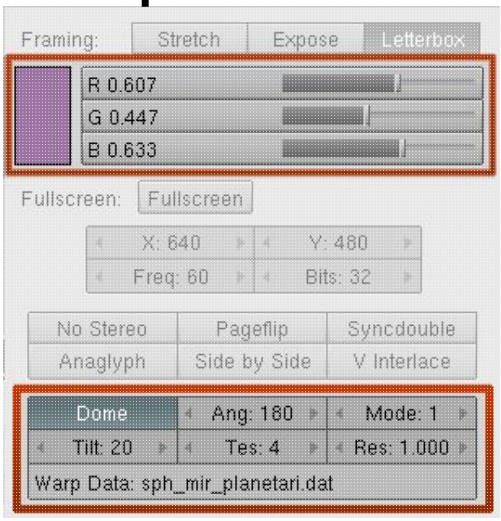
The support has been added to Blender 2.49. It uses a multipass texture algorithm as developed by Paul Bourke. Briefly, that involves rendering the scene 4 times and placing the subsequent images onto a mesh designed especially such that the result, when viewed with an orthographic camera, is a fisheye projection.

Remember to use Blender in fullscreen mode to get the maximum out of your projector.

To accomplish that launch Blender with the command-line argument -W. Also to get away of the top menu on Blender try to join all windows (buttons, 3dview, text, ...) in a single one. Otherwise if you only maximize it (Ctrl+Up) you can't get the whole screen free to run your game (the top bar menu takes about 20 pixels).

*** Dalai Felinto** was sponsored by **SAT -** Society for Arts and Technology within the **SAT Metalab** immersion research program.

GUI options:



- **Background color** (RGB) - the color used as background. Strong colors make calibration/tuning easy.
- **Angle** (from 90 to 250) - it works only in the fisheye mode (1)
- **Mode** (to be extended to implement other modes. So far: 1 (fisheye), 2 (enviroment map) and 3(spherical panoramic)
- **Tilt** (for tilted domes - it rotates the camera in the horizontal axis)
- **Tesselation** (4 is the default. This is the tesselation level of the mesh)
- **Resolution** (1.0 is full resolution. This is the relative size of the render buffer. Decrease it to get speed)
- **Warp Data** (warp mesh data file)

Dome Modes

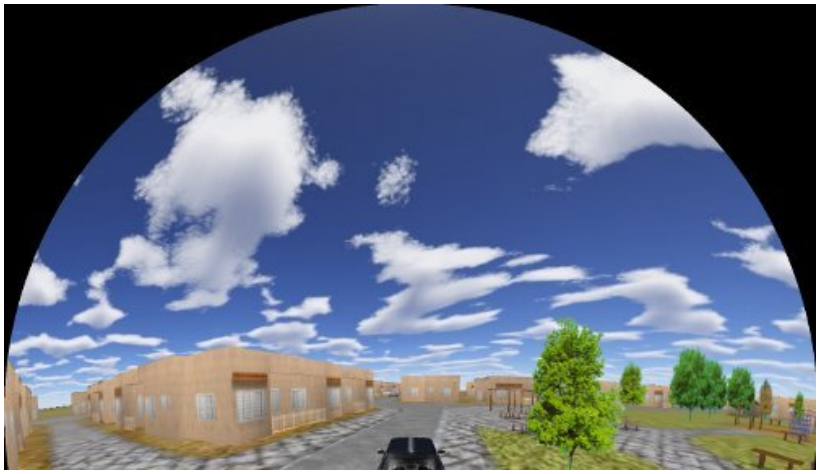
1. Fisheye Dome
2. Front-Truncated Dome
3. Rear-Truncated Dome
4. Environment Map
5. Full Spherical Panoramic

1. Full Dome Mode



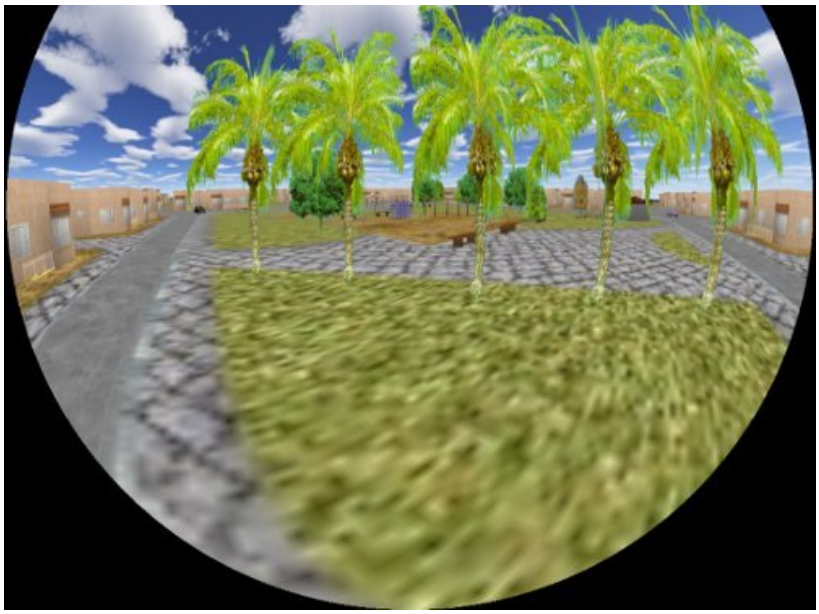
- An Orthogonal Fisheye view from 90Âº to 250Âº degrees.
- From 90Âº to 180Âº we are using 4 renders.
- From 181Âº to 250Âº we are using 5 renders.

2. Front-truncated Dome Mode



- Designed for truncated domes, this mode aligns the fisheye image with the top of the window while touching the sides.
- The Field of view goes from 90Âº to 250Âº degrees.
- From 90Âº to 180Âº we are using 4 renders.
- From 181Âº to 250Âº we are using 5 renders.

3. Rear-truncated Dome Mode



- Designed for truncated domes, this mode aligns the fisheye image with the bottom of the window while touching the sides.
- The Field of view goes from 90Âº to 250Âº degrees.
- From 90Âº to 180Âº we are using 4 renders.
- From 181Âº to 250Âº we are using 5 renders.

4.Environment Map Mode



- An enviroment Map mode. It can be used for pre-generate animated images for CubeMaps.
- We are using 6 renders for that.
- The order of the images follows Blender internal EnvMap file format:
 - ♦ first line: right, back, left
 - ♦ second line: bottom, top, front

5. 360Âº Spherical Panoramic



- A full spherical panoramic mode.
- We are using 6 cameras here.
- The bottom and top start to get precision with **Definition** set to 5 or more.

Warp Data Mesh

> Many projection environments require images that are not simple perspective projections that are the norm for flat screen displays. Examples include geometry correction for cylindrical displays and some new methods of projecting into planetarium domes or upright domes intended for VR.

Paul Bourke



In order to produce that images, we are using a specific file format.

File template::

```
mode
width height
n0_x n0_y n0_u n0_v n0_i
n1_x n1_y n1_u n1_v n1_i
n2_x n1_y n2_u n2_v n2_i
n3_x n3_y n3_u n3_v n3_i
(...)
```

First line is the image type the mesh is support to be applied to: **2** = **rectangular**, **1** = **radial**
Next line has the mesh dimensions in pixels
Rest of the lines are the nodes of the mesh.

Each line is compund of **x y u v i**
(x,y) are the normalised screen coordinates
(u,v) texture coordinates
i a multiplicative intensity factor

x varies from -screen aspect to screen aspect
y varies from -1 to 1

u and v vary from 0 to 1
i ranges from 0 to 1, if negative don't draw that mesh node

- You need to create the file and add it to the Text Editor in order to select it as your Warp Mesh data file.

Example files: Spherical Mirror Dome 4x3, Truncated Dome 4x3, Sample Fullscreen File 4x3, Sample Fullbuffer File 4x3.

Important: the viewport is calculated using the ratio of canvas width by canvas height. Therefore different screen sizes will require different warp mesh files. Also in order to get the correct ratio of your projector you need to use Blender in Fullscreen mode.

Physics Engine

Mode: All Modes

Panel: Shading context â World sub-context â Mist/Stars/Physics

Hotkey: F8

Description

This part of the World determines the type of physical rules that govern the game engine scene, and the gravity of the situation. Based on the physics engine selected, in physics simulations in the game engine, Blender will automatically move Actors in the downward (-Z) direction. After you arrange the actors and they move as you wish, you can then bake this computed motion into fixed Ipo curves (see Logic actors for more info).



Options

Physics
: The type of physics engine to use:

Bullet
: The default physics engine, in active development. It handles movement and collision detection. The things that collide transfer momentum to the collided object.

Sumo (deprecated)
: An older physics engine, used in past versions of Blenderâ s game engine. Use this for compatibility with (very) old files.

None
: No physics in use. Things are not affected by gravity and can fly about in a virtual space. Objects in motion stay in that motion.

Gravity
: The gravitational acceleration, in units of meters per squared second ($m.s^{-2}$), of this world. Each object that is an actor has a mass and size slider (see dynamics materials section). In conjunction with the frame rate (see Render section), Blender uses this info to calculate how fast the object should accelerate downward.

Bullet physics Python API

Bullet Physics provides collision detection and rigid body dynamics for the Blender Game Engine. It takes some settings from Blender that previously were designed for the former collision detection system (called Sumo).

However, new features donâ← t have an user interface yet, so Python can be used to fill the gap for now.

Features:

- Vehicle simulation.
- Rigid body constraints: hinge and point to point (ball socket).
- Access to internal physics settings, like deactivation time, debugging features.

Easiest is to look at the Bullet physics demos, how to use them. More information can be found here.

Python script example:

```
import PhysicsConstraints
print dir(PhysicsConstraints)
```

The VideoTexture module

The `VideoTexture` module allows you to manipulate textures during the game. Several sources for texture are possible: video files, image files, video capture, memory buffer, camera render or a mix of that. The video and image files can be loaded from the internet using an URL instead of a file name. In addition, you can apply filters on the images before sending them to the GPU, allowing video effect: blue screen, color band, gray, normal map. `VideoTexture` uses FFmpeg to load images and videos. All the formats and codecs that FFmpeg supports are supported by `VideoTexture`, including but not limited to:

- AVI
- Ogg
- Xvid
- Theora
- dv1394 camera
- video4linux capture card (this includes many webcams)
- videoForWindows capture card (this includes many webcams)
- JPG

How it works

The principle is simple: first you identify an existing texture by object and name, then you create a new texture with dynamic content and swap the two textures in the GPU. The GE is not aware of the substitution and continues to display the object as always, except that you are now in control of the texture. At the end, the new texture is deleted and the old texture restored.

The present page is a guide to the `VideoTexture` module with simple examples.

Game preparation

Before you can use the thing `VideoTexture` module, you must have objects with textures applied appropriately.

Imagine you want to have a television showing live broadcast programs in the game. You will create a television object and UV-apply a different texture at the place of the screen, for example â← `tv.png`â← . What this texture looks like is not important; probably you want to make it dark grey to simulate power-off state. When the television must be turned on, you create a dynamic texture from a video capture card and use it instead of `tv.png`: the TV screen will come to life.

You have two ways to define textures that `VideoTexture` can grab:

1. Simple UV texture.
2. Blender material with image texture channel.

Because `VideoTexture` works at texture level, it is compatible with all GE fancy texturing features: GLSL, multi-texture, custom shaders, etc.

First example

Letâ← s assume that we have a game object with one or more faces assigned to a material/image on which we want to display a video.

The first step is to create a `Texture` object. We will do it in a script that runs once. It can be at the start of the game, the video is only played when you refresh the texture; weâ← ll come to that later. The script is normally attached to the object on which we want to display the video so that we can easily retrieve the object reference:

```
import VideoTexture

 contr = GameLogic.getCurrentController()
 obj = contr.owner

 if not hasattr(GameLogic, 'video'):
```

The check on â← `video`â← attribute is just a trick to make sure we create the texture only once.

Find material

```
     matID = VideoTexture.materialID(obj, 'IMvideo.png')
```

`VideoTexture.materialID()` is a handy function to retrieve the object material that is using `video.png` as texture. This method will work with Blender material and UV texture. In case of UV texture, it grabs the internal material corresponding to the faces that are assigned to this texture. In case of Blender material, it grabs the material that has an image texture channel matching the name as first channel.

The â← `IM`â← prefix indicates that weâ← re searching for a texture name but we can also search for a material by giving the â← `MA`â← prefix. For example, if we want to find the material called `VideoMat` on this object, the code becomes:

```
     matID = VideoTexture.materialID(obj, 'MAVideoMat')
```

Create texture

`VideoTexture.Texture` is the class that creates the `Texture` object that loads the dynamic texture on the GPU. The constructor takes one mandatory and three optional arguments:

`gameObj`
: The game object.

`materialID`
: Material index as returned by `VideoTexture.materialID()`, 0 = first material by default.

`textureID`
: Texture index in case of multi-texture channel, 0 = first channel by default.
In case of UV texture, this parameter should always be 0.

`textureObj`
: Reference to another `Texture` object of which we want to reuse the texture.

If we use this argument, we should not create any source on this texture and there is no need to refresh it either: the other `Texture` object will provide the texture for both materials/textures.

```
GameLogic.video = VideoTexture.Texture(obj, matID)
```

Make texture persistent

Note that we have assigned the object to a `GameLogic` â← `video`â← attribute that we create for the occasion. The reason is that the `Texture` object must be persistent across the game scripts. A local variable would be deleted at the end of the script and the GPU texture deleted at the same time. `GameLogic` module object is a handy place to store persistent objects.

Create a source

Now we have a `Texture` object but it canâ← t do anything because it does not have any source. We must create a source object from one of the possible sources available in `VideoTexture`:

`VideoFFmpeg`
: Moving pictures.
Video file, video capture, video streaming.

`ImageFFmpeg`
: Still pictures.
Image file, image on web.

`ImageBuff`
: Image from application memory.
For computer generated images, drawing applications.

`ImageViewport`
: Part or whole of the viewport (=rendering of the active camera displayed on screen).

`ImageRender`
: Render of a non active camera.

`ImageMix`
: A mix of 2 or more of the above sources.

In this example we use a simple video file as source. The `VideoFFmpeg` constructor takes a file name as argument. To avoid any confusion with the location of the file, we will use `GameLogic.expandPath()` to build an absolute file name, assuming the video file is in the same directory as the blend file:

```
movie = GameLogic.expandPath('//trailer_400p.ogg')
GameLogic.video.source = VideoTexture.VideoFFmpeg(movie)
```

We create the video source object and assign it to the `Texture` object `source` attribute to set the source and make it persistent: as the `Texture` object is persistent, the source object will also be persistent.

Note that we can change the `Texture` source at any time. Suppose we want to switch between two movies during the game. We can do the following:

```
GameLogic.mySources[0] = VideoTexture.VideoFFmpeg('movie1.avi')
GameLogic.mySources[1] = VideoTexture.VideoFFmpeg('movie2.avi')
```

And then assign (and reassign) the source during the game:

```
GameLogic.video.source = GameLogic.mySources[movieSel]
```

Setup the source

The `VideoFFmpeg` source has several attributes to control the movie playback:

`range`
: [start,stop] (*floats*).
Set the start and stop time of the video playback, expressed in seconds from beginning. By default the entire video.

`repeat`
: (*integer*).
Number of video replay, -1 for infinite.

`framerate`
: (*float*).
Relative frame rate, <1.0 for slow, >1.0 for fast.

`scale`
: (*bool*).
Set to True to activate fast nearest neighbour scaling algorithm.
Texture width and height must be a power of 2. If the video picture size is not a power of 2, rescaling is required. By default `VideoTexture` uses the precise but slow `gluScaleImage()` function.
Best is to rescale the video offline so that no scaling is necessary at runtime!

`flip`
: (*bool*).
Set to True if the image must be vertically flipped.
FFmpeg always delivers the image upside down, so this attribute is set to True by default.

`filter`
: Set additional filter on the video before sending to GPU.
Assign to one of `VideoTexture` filter object. By default the image is send unchanged to the GPU.
If an alpha channel is present in the video, it is automatically loaded and sent to the GPU as well.

We will simply set the `scale` attribute to True because the `gluScaleImage()` is really too slow for real time video. In case the video dimensions are already a power of 2, it has no effect.

```
GameLogic.video.source.scale = True
```

Play the video

We are now ready to play the video:

```
GameLogic.video.source.play()
```

Video playback is not a background process: it happens only when we refresh the texture. So we must have another script that runs on every frame and calls the `refresh()` method of the `Texture` object:

```
if hasattr(GameLogic, 'video'):
     GameLogic.video.refresh(True)
```

If the video source is stopped, `refresh()` has no effect. The argument of `refresh()` is a flag that indicates if the texture should be recalculated on next refresh. For video playback, you definitively want to set it to True.

Checking video status

Video source classes (such as VideoFFMpeg) have an attribute `status`. If video is playing, its value is 2, if it's stopped, it's 3. So in our example:

```
if GameLogic.video.source.status == 3:
     #video has stopped
```

Advanced work flow

True argument in `Texture.refresh()` method simply invalidates the image buffer after sending it to the GPU so that on next frame, a new image will be loaded from the source. It has the side effect of making the image unavailable to Python. You can also do it manually by calling the `refresh()` method of the source directly.

Here are some possible advanced work flow:

- Use the image buffer in python (doesn't effect the Texture):

```
GameLogic.video.refresh(False)
image = GameLogic.video.source.image
# image is a binary string buffer of row major RGBA pixels
# ... use image
# invalidates it for next frame
GameLogic.video.source.refresh()
```

- Load image from source for python processing wihtout download to GPU:

```
# note that we don't even call refresh on the Texture
# we could also just create a source object without a Texture object

image = GameLogic.video.source.image
# ... use image
GameLogic.video.source.refresh()
```

Download the demo

You can find the demo here.

Advanced demos

Here is a demo that demonstrates the use of two videos alternatively on the same texture. Note that it requires an additional video file which is the elephant dream teaser. You can replace with another other file that you want to run the demo.

Here is a demo that demonstrates the use of the `ImageMix` source. `ImageMix` is a source that needs sources, which can be any other `Texture` source, like `VideoFFmpeg`, `ImageFFmpeg` or `ImageRender`. You set them with `setSource()` and their relative weight with `setWeight()`. Pay attention that the weight is a short number between 0 and 255, and that the sum of all weights should be 255. `ImageMix` makes a mix of all the sources according to their weights. The sources must all have the same image size (after reduction to the nearest power of 2 dimension). If they donâ ← t, you get a Python error on the console.

Features

Brief incomplete list of features of the Game Engine:

- Platforms: GNU/Linux, Windows, Macâ←¦
- Logicbrick system.
- Python scripting.
- Supports GLSL shader 2.0 (vertex and fragments programs via Python or build in).
 - ♦ Normalmap.
 - ♦ Parallaxmap.
 - ♦ Specularmap.
 - ♦ Colormap.
 - ♦ Detailmap.
 - ♦ Environment reflect cube.
 - ♦ Environment reflect sphere.
 - ♦ Vertex color.
 - ♦ Phong shader.
 - ♦ Dynamic shadow (soft and stencil shadow - not yet).
 - ♦ GL particle system (not yet).
- Physics: Bullet and Sumo.
- Sound: wav (ogg,mp3 not yet).
- Textures: tga+alpha, png+alpha, jpg.
- Bitmap fonts.
- GUI (graphical user interface - not yet).
- Mesh deform with bones.
- Terrain engine (not yet).
- Lod (level of detail - via Python).
- Culling system: sphere, box, frustum (via Python).
- Portals (not yet).
- Network (via Python).
- Multi uv coordinate (via Python).
- Mipmap.
- Lightmap/shadowmap baking.

External Resources

External documentation

- Gamekit 2.0
- Original Blender GameKit
- Using Material GLSL Shaders
- External tutorials

Community

- Blenderartists Game Engine Support and Discussion forum
- #gameblender Chat IRC channel

Development

- Game Engine Development (wiki pages)
- Game engine bug tracker

Game Development tools and utilities

- GIMP - â← GNU Image Manipulation Programâ← , on GNU/Linux and Windows
- IconArt - make custom icons for your game.exe
- Setup2go and Installer2go - installer creators
- Inno Setup - a very good free installer
- BZoo - a powerful network game template
- Game Engine Competition
- WSAG - A Solid Multi Player Online Network

Python

- Python notes
- BGE Python API Docs (2.49)